arXiv:1403.4630v1 [stat.ME] 18 Mar 2014

Penalising model component complexity: A principled, practical approach to constructing priors

Thiago G. Martins, Daniel P. Simpson, Andrea Riebler & Håvard Rue
Department of Mathematical Sciences
NTNU, Norway

Sigrunn H. Sørbye
Department of Mathematics and Statistics
UiT The Arctic University of Norway

March 20, 2014

Abstract

The issue of setting prior distributions on model parameters, or to attribute uncertainty for model parameters, is a difficult issue in applied Bayesian statistics. Although the prior distribution should ideally encode the users' prior knowledge about the parameters, this level of knowledge transfer seems to be unattainable in practice and applied statisticians are forced to search for a "default" prior. Despite the development of objective priors, which are only available explicitly for a small number of highly restricted model classes, the applied statistician has few practical guidelines to follow when choosing the priors. An easy way out of this dilemma is to re-use prior choices of others, with an appropriate reference.

In this paper, we introduce a new concept for constructing prior distributions. We exploit the natural nested structure inherent to many model components, which defines the model component to be a flexible extension of a base model. Proper priors are defined to penalise the complexity induced by deviating from the simpler base model and are formulated after the input of a user-defined *scaling* parameter for that model component.

These priors are invariant to reparameterisations, have a natural connection to Jeffreys' priors, are designed to support Occam's razor and seem to have excellent robustness properties, all which are highly desirable and allow us to use this approach to define default prior distributions.

KEYWORDS: Bayesian theory, Disease mapping, Information geometry, Interpretable prior distributions, Jeffreys' prior, Prior on Toeplitz matrices, Prior on correlation matrices

AMS SUBJECT CLASSIFICATION:: Primary 62F15; secondary 62H99

ADDRESS FOR CORRESPONDENCE: H. Rue, Department of Mathematical Sciences, The Norwegian University for Science and Technology, N-7491 Trondheim, Norway. Email: hrue@math.ntnu.no, WWW-address: http://www.math.ntnu.no/~hrue/

1 Introduction

The core of Bayesian statistics is to use Bayes' theorem to update prior beliefs about events. The choice of prior distributions, except for simple cases, is however often problematic (Kass and Wasserman, 1996). Prior elicitation from an expert is difficult in practice, but parameters in a parametric family can sometimes be elicited implicitly from interviewing experts to construct an informative prior (Kadane et al., 1980; O'Hagan et al., 2006; Albert et al., 2012) also in higher dimension (Cooke, 1991; Kurowicka and Cooke, 2006). However, it is usually much more difficult to prescribe sensible priors on parameters further down the hierarchy, such as precision or correlation parameters. Martyn Plummer, the author of `JAGS` software for Bayesian inference (Plummer, 2003) goes so far as to say[1] "*[...] nobody can express an informative prior in terms of the precision [...]*".

On the other end of the hunt for the holy grail, "objective" priors (Bernardo, 1979; Berger, 2006; Berger et al., 2009a; Ghosh, 2011) are data-dependent and are not uniformly accepted among Bayesians on philosophical grounds; see for example discussion contributions to Berger (2006) and Goldstein (2006), but these results can still be useful in practice. Jeffreys' non-informative priors–and their extension "reference priors" (Berger et al., 2009b)—remove this philosophical difficulty, but great care is needed to ensure that the resulting posteriors are proper. If chosen appropriately, the use of non-informative priors will lead to appropriate parameters estimates as demonstrated on several applications by Kamary (2014). However, these priors are model dependent and notoriously difficult to derive except for simple cases. Furthermore, when the likelihood changes or another component is added to the linear predictor, the entire prior must be recomputed. This does not mesh well with the practice of applied statistics, in which a "building block" approach is commonly used and several slightly different models will be fitted simultaneously to the same dataset. Furthermore, the use of improper priors wreaks havoc on marginal likelihood-based model selection procedures.

Between these two extremes lies the realm of "weakly informative" priors (Gelman, 2006; Gelman et al., 2008). These priors are constructed by recognised that while you usually do not have strong prior information about the value of a parameter, it is rare to be completely ignorant. For example, when estimating the height and weight of a person, it is sensible to select a prior that gives mass neither to people who are five metres tall, nor to those who only weight two kilograms. This use of weak prior knowledge is sufficient to regularise the extreme inferences that are obtained using maximum likelihood or non-informative priors.

Despite the challenge of choosing priors, more and more researchers and practitioners apply Bayesian methods to analyse their data and make inference. The most common approach to prior selection is, unfortunately, to use a prior from the literature. This is rarely a good idea, especially when this prior was originally chosen based on a specific property of the dataset in question. Another source of potentially disastrous priors are those derived for mathematical or computational convenience. One extreme example of this is the enduring popularity of the $\Gamma(\epsilon, \epsilon)$ prior, with a small ϵ, for inverse variance (precision) parameters, which is the "default" choice in the WinBugs (Spiegelhalter et al., 1995) example manuals, although it is well known to be a choice with severe bi-effects; see the discussion in Fong et al. (2010) and Hodges (2013). In essence, the prevailing practice is to try to avoid the problem by hoping that the data will dominate the prior, or to defer to the choices of others. This is not how Bayesian statistics is meant to be.

The problem of constructing sensible default priors on model components becomes especially pressing when developing general software for Bayesian computation. As developers of the `R-INLA` (see `http://www.r-inla.org/`) package, which performs approximate Bayesian inference for latent Gaussian models (Rue et al., 2009; Lindgren et al., 2011; Martins et al., 2013), we are left with two unpalatable choices. We could force the user of `R-INLA` to explicitly define prior distributions for all parameters in the model. Arguably, this is the correct thing to do, however, the aforementioned sea of confusion around how to properly prescribe priors makes this impractical. The second option is to provide default reference or objective priors for all user-defined models, which is also highly impractical due to their model-dependent nature. The current approach taken in `R-INLA` is, therefore, to provide a set of fixed default priors, selected partly out of the blue without any underlying principles. The novel approach presented in this paper aims at improving this

[1] `http://martynplummer.wordpress.com/2012/09/02/stan/`

suboptimal practice.

Throughout this paper, we will consider an additive approach to Bayesian model building, in which a model is decomposed into discrete building blocks such as i.i.d. Gaussian random effects, spline models or spatial effects. Most model components can be naturally regarded as a flexible version of a *base model*. For example the i.i.d. Gaussian random effect can be seen as a generalisation of the base model which puts all prior mass on zero, i.e. there is no random effect.

In the presence of a base model, constructing priors becomes somewhat simpler. We use, as a starting point for our considerations, the following loose definition.

Definition 1. *A prior* $\pi(\theta)$ overfits *if it the density of prior is zero at the base model. This implies vanishing prior mass in the neighbourhood of the base model which means that the posterior cannot concentrate on the true model.*

Hence the prior will drag the posterior towards a more complex posterior. In this paper, we derive priors that nstead shrink towards the base model. In Sections 4 and 5, we show that this property really matters in practice, especially in the presence of semiparametric effects. In order to successfully apply Definition 1, it is necessary to carefully consider what a neighbourhood of the base model looks like. In Section 2 we show that it is not enough to consider a fixed parameterisation, but rather it is necessary to consider the information geometry of the parameter space.

The goal of this paper is to construct *penalised complexity priors* (PC priors) that allow the user to define a notion of *scale* for each model component, which is then used to *define* a reasonable prior using our new principled approach. Thus, we are able to derive, for all model components, unified prior specifications with a clear meaning and interpretation. The principles are designed so that desirable properties follow automatically: invariance regarding reparameterisations, connection to Jeffreys' prior, support of Occam's razor principle, and empirical robustness to the choice of scaling parameter. Our approach is not restricted to any specific computational method as it is a principled approach to prior construction and therefore relevant to any application involving Bayesian analysis.

As a simple example, consider the following Poisson regression model

$$y_i \mid \ldots \sim \text{Poisson}\left(\exp(\mu + \beta z_i + u_i)\right)$$

where μ is the intercept, β is the effect of covariate $\boldsymbol{z}$ and $\boldsymbol{u} \sim N(\boldsymbol{0}, \tau^{-1}\boldsymbol{I})$ is a Gaussian random effect. The only hyper-parameter in this model is the precision of $\boldsymbol{u}$, called τ, to which we need to assign a prior distribution. Even in this simple case it is difficult to define a sound prior. Applying the new principled approach developed below, the current model is regarded as a flexible extension of a base model $\tau = \infty$, i.e. no random effect. The user must provide a scale U, with the interpretation

$$\text{Prob}(1/\sqrt{\tau} > U) = \alpha$$

where α is a predefined small probability, such as $\alpha = 1\%$. This value encodes information about the scale of $\boldsymbol{u}$ so that the event $|u_i| > 4U$ is somewhat unlikely. Using our new approach, we can derive the following prior distribution for τ,

$$\pi(\tau) = \frac{\lambda}{2}\tau^{-3/2}\exp\left(-\lambda\tau^{-1/2}\right), \qquad \tau > 0$$

which is a type-2 Gumbel distribution, where λ is a function of U and α. This prior also corresponds to an exponential prior with rate λ for the standard deviation. A curious property, is that the expected value of τ does not exist. This is a consequence of Occam's razor that is built into our approach. The random effect model is interpreted as a flexible extension of a base model, in this case the constant zero (i.e. no random effect), hence Definition 1 states that $\tau = \infty$ must be attainable in the prior in order to avoid overfitting.

The plan of the paper is as follows. In Section 2 we will define our principled approach to design priors and discuss its connection to the Jeffreys' prior. In Section 3 we study prior choices for the degrees of freedom parameter in the Student-t distribution and perform an extensive simulation study to investigate properties of the new prior. In Section 4 we consider a Cox proportional hazard model for the analysis of recurrent

time to event data. We discuss prior choices for the precision and the correlation parameter for a time-dependent frailty model component and priors for the precision parameter in the spline model component of the log baseline hazard. In Section 5 we discuss the BYM-model for disease mapping with a possible smooth effect of an ecological covariate, and we suggest a new reparameterisation of the model in order to facilitate improved control and interpretation. Section 6 extends the method to multivariate parameters and we derive new, principled priors for correlation matrices and the coefficients of auto-regressive processes. We end with a discussion in Section 7. The Appendix will host some technical details.

2 Penalised complexity priors

In this section we will outline our approach to defining PC priors for a univariate parameter, postponing the extensions to the multivariate case. We will also discuss the connection to Jeffreys' prior.

2.1 Flexible extensions of base models

Let $\pi(\boldsymbol{x}|\xi)$ be the density for the model component of interest where the task is to construct a reasonable prior for the parameter ξ. For ease of notation and without loss of generality, we will assume that $\xi \geq 0$, and the base model corresponds to $\xi = 0$, that is that $\pi(\boldsymbol{x}|\xi = 0)$ is a natural and simpler base model for $\pi(\boldsymbol{x}|\xi)$. This allows us to interpret $\pi(\boldsymbol{x}|\xi)$ as a *flexible extension* of the base model, where increasing values of ξ imply increasing deviations from the base model. A few simple examples will fix the ideas.

Student-t response Let $\boldsymbol{x}|\xi$ be independent Student-t densities with unit variance and ν degrees-of-freedom where $\nu = 1/\xi$. The base-model is the Gaussian with zero mean and precision matrix $\boldsymbol{I}$. The Student-t can then be interpreted as a flexible extension of the Gaussian which is appropriate for robust likelihood and robust random effects models.

Gaussian random effects Let $\boldsymbol{x}|\xi$ be multivariate Gaussians with zero mean and precision matrix $\tau\boldsymbol{I}$ where $\tau = 1/\xi$. Here, the base-model puts all the mass at zero, which is appropriate for random effects where the natural reference is absence of these effects. In the multivariate case and conditional on τ, we can allow for correlation among the components where the uncorrelated case is the base-model.

Spline model Let $\boldsymbol{x}|\xi$ represent a spline model with smoothing parameter $\tau = 1/\xi$. The base model is the infinite smooth spline which can be a constant or a straight line, depending on the order of the spline or in general its null-space. This interpretation is natural when the spline model is used as a flexible extension to a constant (see Section 4.2) or in generalised additive models, which can (and should) be viewed as a flexible extension of a generalised linear model.

Time dependence Let $\boldsymbol{x}|\xi$ denote an auto-regressive model of order 1, unit variance and lag-one correlation $\rho = \xi$. Depending on the application, then either $\rho = 0$ (no dependence in time) or $\rho = 1$ (no changes in time) is the appropriate base-model.

2.2 A principled definition of the PC prior

We will now state and discuss our principles for constructing a prior distribution for ξ.

Occam's razor is the principle of parsimony, for which simpler model formulations should be preferred until there is enough support for a more complex model. We want to incorporate this principle into the prior. Our simpler model is the base-model hence we want the prior to penalise deviations from it. From the prior alone we should prefer the simpler model and the prior should be decaying as a function of a measure of the increased complexity between the more flexible model and the base model.

Measure of complexity We will use the Kullback-Leibler divergence (KLD) to measure the increased complexity (Kullback and Leibler, 1951). Between densities f and g, the KLD is defined by

$$\text{KLD}\left(f\|g\right) = \int f(\boldsymbol{x}) \log\left(\frac{f(\boldsymbol{x})}{g(\boldsymbol{x})}\right) \, d\boldsymbol{x}.$$

KLD is a measure of the information lost when the base model g is used to approximate the more flexible model f. The asymmetry in this definition is vital to the success of this method: Occam's razor is naturally non-symmetric and, therefore, our measure of complexity must also be non-symmetric. In order to use the KLD, it is, however, necessarily to transform it onto a physically interpretable "distance" scale. Throughout this paper, we will use the (unidirectional) measure $d(f||g) = \sqrt{2\text{KLD}\left(f\|g\right)}$, to define the distance between the two models. (We will use this phrase even though it is not a metric.) Hence, we consider d to be a measure of complexity of the model f when compared to model g. Here, the factor "2" is introduced for convenience and the square root deals with the power of two that is naturally associated with the KLD (see the simple example below).

Constant rate penalisation Penalising the deviation from the base model parametrised with the distance d, we use a constant decay-rate r, so that the prior satisfies

$$\frac{\pi_d(d+\delta)}{\pi_d(d)} = r^\delta, \qquad d, \delta \geq 0$$

for some constant $0 < r < 1$. This will ensure that the relative prior change by an extra δ does not depend on d, which is a reasonable choice without extra knowledge. Further, the mode of the prior is at $d = 0$, i.e. the base model. The constant rate penalisation assumptions imply an exponential prior on the distance scale, $\pi(d) = \lambda \exp(-\lambda d)$, for $r = \exp(-\lambda)$. This corresponds to the following prior on the original space

$$\pi(\xi) = \lambda \mathrm{e}^{-\lambda d(\xi)} \left|\frac{\partial d(\xi)}{\partial \xi}\right|. \tag{1}$$

In some cases d is upper bounded and we use a truncated exponential as the prior for d.

User-defined scaling The final principle needed to completely define a PC prior is that the *user* has an idea of a sensible size for the parameter or a property of the correponding model component. This is similar to the principle behind weakly informative priors. In this context, we can select λ by controlling the prior mass in the tail. This condition is of the form

$$\Pr(Q(\xi) > U) = \alpha,$$

where $Q(\xi)$ is an interpretable transformation of the flexibility parameter, U is a "sensible", user-defined upper bound that specifies what we think of as a "tail event", and α is the weight we put on this event. This condition allows the user to prescribe how informative the resulting PC prior is.

The following simple example will illustrate the ideas. Let $\mathcal{N}(x; \xi, 1)$, $\xi > 0$, be the Gaussian flexible model with mean ξ and precision 1, and where $\xi = 0$ is the base model. The KLD between the two Gaussians is $\xi^2/2$, meaning that $d(\xi) = \sqrt{2\text{KLD}(\xi)} = \xi$, and the trivial change-of-variable exercise from $d(\xi)$ to ξ results in an exponential prior for ξ,

$$\pi(\xi) = \lambda \exp\left(-\lambda\xi\right), \qquad \xi \geq 0$$

for some $\lambda > 0$. The parameter λ has to be decided by the user, providing some knowledge on the scale of ξ, for example $\text{Prob}(\xi > 1) = 1\%$

This procedure is invariant to reparameterisation, since the prior is defined on the distance d, which is then transformed to the corresponding prior for ξ. This is a major advantage of PC priors, since we can construct the prior without taking the specific parameterisation into account.

The PC prior is not defined for a fixed effect, since the coefficient is not a parameter in a model component, but defines a model component by itself.

2.3 Behaviour near the base model

To understand this new construction better, we can study what happens near $\xi = 0$ when it is a regular point, using the connection between KLD and the Fisher information metric. Let $I(\xi)$ be the Fisher information at ξ. Using the well known asymptotic expansion

$$\text{KLD}\left(\pi(\boldsymbol{x} \mid \xi) \| \pi(\boldsymbol{x} \mid \xi = 0)\right) = \frac{1}{2} I(0) \xi^2 + \text{higher order terms},$$

a standard expansion reveals that our new prior behaves like

$$\pi(\xi) = I(\xi)^{1/2} \exp\left(-\lambda\, m(\xi)\right) + \text{higher order terms} \tag{2}$$

for $\lambda\xi$ close to zero. Here, $m(\xi)$ is the distance defined by the metric tensor $I(\xi)$, $m(\xi) = \int_0^\xi \sqrt{I(s)} ds$, using tools from information geometry. Close to the base model, the PC prior is a tilted Jeffreys' prior for $\pi(\boldsymbol{x}|\xi)$, where the amount of tilting is determined by the distance on the Riemannian manifold to the base model scaled by the parameter λ. The user-defined parameter λ thus determines the degree of informativeness in the prior.

2.4 PC priors, non-informative priors and divergence-based priors

We stress that we neither aim to prescribe, nor want to use, non-informative or objective priors for our parameter ξ. Firstly, the main idea behind non-informative priors is the absence of prior knowledge. The principles outlined above are specifically built for the case where we have some prior information, namely the existence of a base model. Furthermore, we assume that the user has some idea of a sensible scale for the flexible model component. Therefore, our approach is closer in spirit of weakly informative priors defended by Gelman (2006); Gelman et al. (2008), although, should the user require it, there is no reason that these priors can't be truly informative. Secondly, non-informative/objective priors necessaritly depend on the model structure and have to be recomputed for each new model. For a general-purpose software system like `R-INLA`, `OpenBUGS` (Lunn et al., 2009), `JAGS` (Plummer, 2003), or `STAN` (Stan Development Team, 2014) this is infeasible.

The idea of using the distance between two models as a reasonable scale to think about priors dates back to Jeffreys (1946) pioneering work to obtain priors that are invariant to reparameterisation. Bayarri and García Donato (2008) build on the early ideas of Jeffreys to derive objective priors for computing Bayes factors for Bayesian hypothesis tests; see also Robert et al. (2009, Sec. 6.4). They use divergence measures between the competing models to derive the required proper priors, and call those derived priors divergence-based (DB) priors. Given the prior distribution on the parameter space of a full encompassing model, Consonni and Veronese (2008) used Kullback-Leibler projection, in the context of Bayesian model comparison, to derive suitable prior distributions for candidate submodels.

3 The Student-t case

In this section we will study the Student-t case focusing solely on the degrees of freedom (d.o.f.) parameter ν, keeping the precision fixed. This is an important non-trivial case, since the Student-t distribution is often used to robustify the Gaussian distribution. Inference based on the Gaussian distribution is well known to be vulnerable to model deviation. This can result in a significant degradation of estimation performance (Lange et al., 1989; Pinheiro et al., 2001; Masreliez and Martin, 1977), and using the Student-t distribution can account for deviations caused by heavier tails, see for example Chib et al. (2002) and Jacquier et al. (1994) for applications in the econometric literature.

The base model for the Student-t distribution is the Gaussian, which occurs when $\nu = \infty$. In order to maintain the interpretability of the precision parameter of the distribution when $\nu < \infty$, we will use a standardised version of the Student-t with unit precision for all $\nu > 2$. This follows the advice from Cox and Reid (1987), promoting more orthogonal parameters and will ease the (later) joint prior specification of the precision and the d.o.f. Our interpretation of the Occam's razor principle implies that the mode of $\pi_d(d)$ must

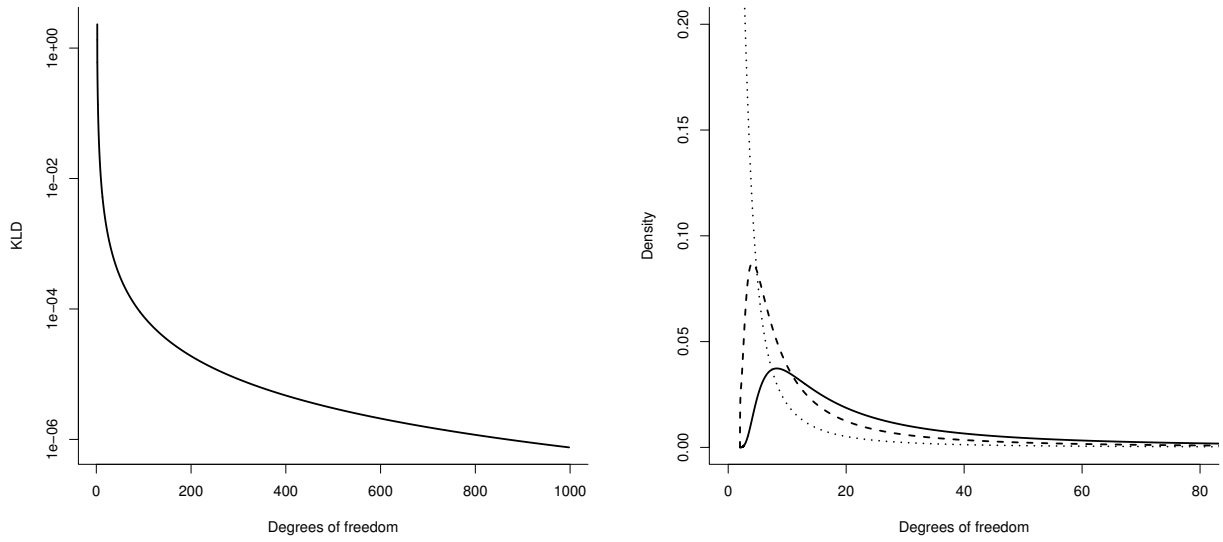

(a) (b)

Figure 1: Panel (a) shows the KLD between the Student-t and the Gaussian distribution (with fixed precision) as a function of the degrees of freedom ν. Panel (b) shows the new priors for ν with $\nu^\star = 10$ and $\alpha = 0.2$ (solid), $\alpha = 0.5$ (dashed) and $\alpha = 0.8$ (dotted).

be at $d = 0$, corresponding to the Gaussian distribution. It turns out that any proper prior for ν with finite expectation violates this principle and promotes overfitting, as $\pi_d(0) = 0$.

Theorem 1. *Let $\pi_\nu(\nu)$ be a prior for $\nu > 2$ where $E(\nu) < \infty$, then $\pi_d(0) = 0$ and the prior overfits.*

The proof is given in Appendix B.1. The intuition is quite clear; if we want $\nu = \infty$ to be central in the prior, a finite expectation will bound the tail behaviour so that we cannot have the mode (or a non-zero density) at $d = 0$. Furthermore, commonly used priors for the d.o.f. often include the exponential (Geweke, 2006) or the uniform (on a finite interval) (Jacquier et al., 2004), which might be due to its use in the examples of the WinBugs manual. These choices are often justified from mathematical tractability or simplicity, but place zero density mass onto the base model. When the Student-t distribution is used as a robust extension of the Gaussian, these priors simply disregard the Gaussian and in that sense overfit. Notably exceptions are the work of Fonseca et al. (2008) which computed (various forms of) Jeffreys' priors in the case of linear regression models with Student-t errors and Villa and Walker (2014) who provide an objective prior for discrete values of the d.o.f.; Please refer to our discussion on non-informative/objective priors in Section 2.4.

Figure 1 (a) shows the KLD as a function of d.o.f. up to $\nu = 1000$. We compute this numerically for $\nu \leq 9$ and use an accurate asymptotic expansion, given in Appendix C, for $\nu > 9$ with maximum absolute error less than 1.6E–10. We then use spline-interpolation to construct a continuous function for KLD and its inverse. The notion of *scale* required from the user, can be defined in several ways. A simple choice is to provide $(\nu^\star, \alpha)$ such that $\text{Prob}(\nu < \nu^\star) = \alpha$, giving $\lambda = -\log(\alpha)/d(\nu^\star)$ in the prior for d. Here, we fix $\nu^\star = 10$ and use in the rest of this section $\alpha = 0.2, 0.3, 0.4, 0.5, 0.6, 0.7$ and 0.8, to denote an increasing deviance from the Gaussian base model. Figure 1 (b) shows the corresponding priors for ν with $\alpha = 0.2, 0.5$ and 0.8.

To illustrate the importance of Theorem 1, consider the exponential prior for $\nu > 2$ with mean equal to 5, 10 and 20, For these choices, the condition for Theorem 1 applies. Figure 2 (a) displays these seemingly reasonable priors for ν converted to a prior for the distance $d = \sqrt{2\ \text{KLD}}$. Figure 2 (b) displays the corresponding uniform priors from $\nu = 2$ to 20, 50 and 100. As predicted, the density at $\nu \to 2$ is zero for all the six priors. For the exponential priors in panel (a), the mode corresponds to $\nu = 7.0$, 17.0 and 37.0, respectively. In practical terms, this implies that Occam's razor does not apply to this prior, in the sense that the exponential prior simply *defines* that the Student-t distribution is genuinely different from the Gaussian base model and the posterior will shrink towards the mode rather than towards the Gaussian base model. The uniform prior behaves similarly as we increase the upper limit, we put more mass to large d.o.fs and the mode

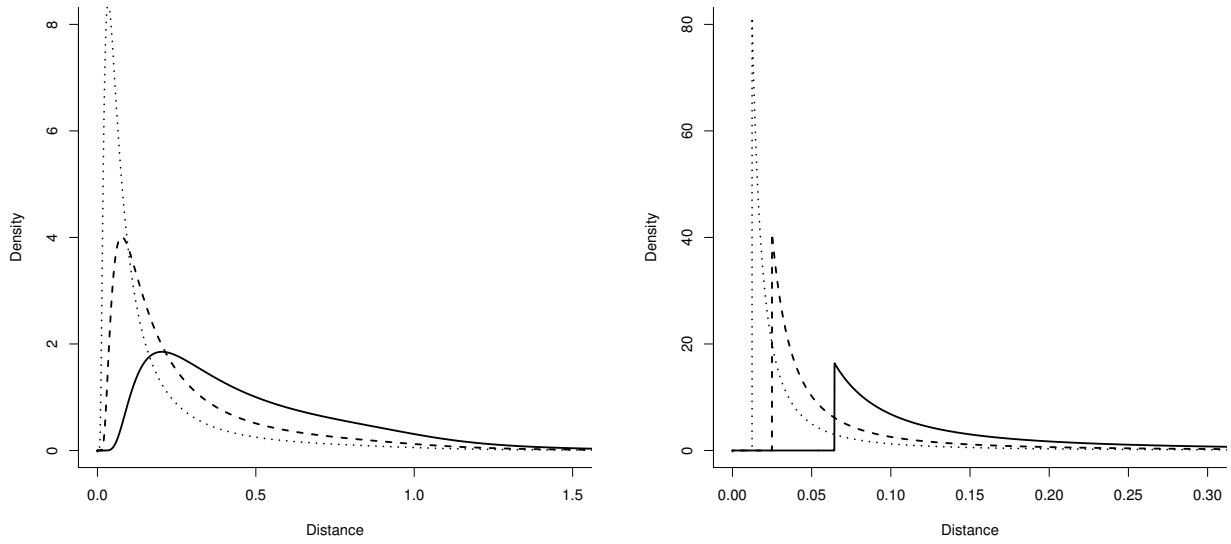

(a) (b)

Figure 2: Panel (a) shows the exponential prior with mean for the d.o.f. of 5 (solid), 10 (dashed) and 20 (dotted) $\nu > 2$ shown on the distance scale. The mode corresponds to $\nu = 7.0$, 17.0 and 37.0, respectively. Panel (b) shows the uniform prior for the d.o.f. from 2, to 20 (solid), 50 (dashed) and 100 (dotted), converted to the distance scale.

moves to left. The finite support implies that the density is zero up the point defined by the upper limit. If the true distribution was Gaussian then we would over-fit the data using any of these priors. This illustrates that the d.o.f. parameterisation is somewhat counter-intuitive and not particular user-friendly. The invariance property of our approach, however, is very useful in such cases, as the prior is defined on properties of the densities and hence is invariant under reparameterisations.

To investigate the properties of the new prior on ν, we will present results from a simulation experiment, where we used the simple model $y_i = \epsilon_i$, $i = 1, \ldots, n$. Similar results are obtained for more involved models (Martins and Rue, 2013). The noise term ϵ is Student-t distributed, with unknown d.o.f. and fixed unit precision.

We simulated data sets for $n = 100, 1000, 10000$. For the d.o.f. we used $\nu = 5, 10, 20, 100$, to study the effect of the priors under different levels of the kurtosis. This gives 12 different scenarios. For each scenario we simulated 1000 different data sets. For each data set we computed the posterior distribution of ν and formed the equal-weight mixture over all the 1000 realisations to approximate the expected behaviour of the posterior distribution over different realisations of the data. The results are shown in Figure 3 using our new prior with $\nu^\star = 10$ and $\alpha = 0.2, 0.3, 0.4, 0.5, 0.6, 0.7$ and 0.8, and the exponential prior with mean 5, 10, 20 and 100.

Each row in Figure 3 displays results for a different d.o.f. while each column displays results for a different sample size n. In each scenario we display the 0.025, 0.5 and 0.975-quantiles of the mixtures of posterior distributions of ν associated to different prior distributions for ν. For example, the first row and the first column in Figure 3 shows credible intervals for each of the priors used in our experiment when $n = 100$ and $\nu = 100$.

The first row in Figure 3 displays the results with $\nu = 100$ which is close to Gaussian observations. The results obtained with the new priors present wide credible intervals covering most scenarios in the presence of few data points but quickly learn about the high d.o.f. in the data as more data is provided. Looking at the results obtained with the exponential priors, we see that the results are highly dependent on the mean in the exponential prior as the results for $n = 100$ change abruptly if the mean is 100 or 10. This difference seems to remain even with $n = 1000$ and $n = 10000$, which could be explained by the fact that the prior still dominates the data. The intervals obtained with the exponential prior for ν look similar for all scenarios, with the exception of the scenarios with low d.o.f. and high sample size, for which there are strong information in

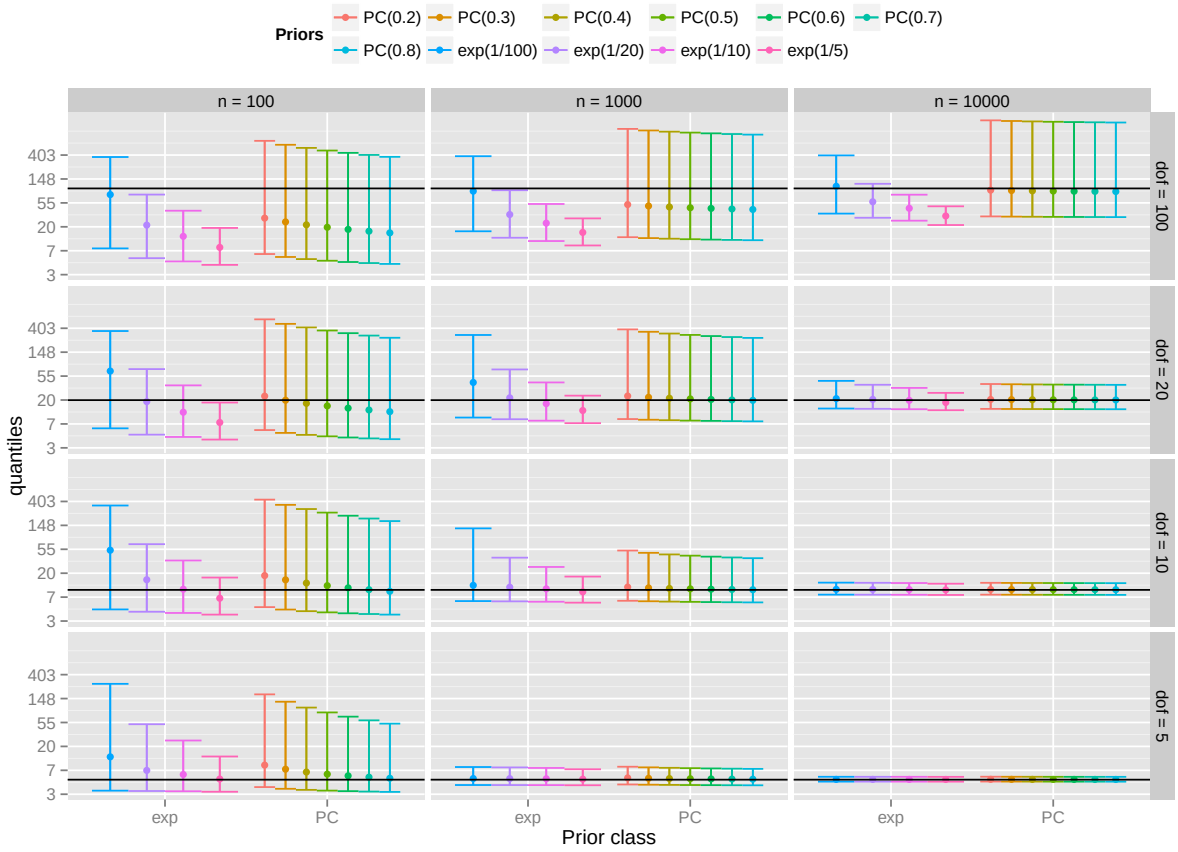

Figure 3: The 0.025-, 0.5- and 0.975-quantile estimates obtained from an equal-weight mixture of posterior distributions of ν obtained by fitting a Student-t errors model with different priors for ν over 1000 datasets, for each of the 16 scenarios with sample sizes $n = 100, 1000, 10000$ and d.o.f. $\nu = 100, 20, 10, 5$. The first seven intervals in each scenario correspond to the new PC prior with $\nu^\star = 10$ and $\alpha = 0.2, 0.3, 0.4, 0.5, 0.6, 0.7$ and 0.8. The last four intervals in each scenario correspond to exponential priors with mean 100, 20, 10, 5, respectively.

the data to dominate the highly informative exponential prior.

If we study Figure 3 column-wise and top-down, we see the difference in the inference obtained with the new priors when we go from high to low d.o.f. Firstly, we note that the performance of the PC priors are barely effected by the change in α and they perform well for all sample sizes. For the exponential priors, we basically see no difference in inference for ν when $n = 100$, comparing the near Gaussian scenario ($\nu = 100$) with the strongly heavy tailed one ($\nu = 5$). The implication is that the results will be much more influenced by the choice of the mean in the exponential prior than by the d.o.f. in the data. Similarly, the exponential priors continue to be highly informative even for large sample sizes. This informative behaviour can be seen in particular in the first row ($\nu = \infty$), and is consistent with Definition 1: exponential priors force overfitting when the data comes from the base model.

This example sheds light on the consistency issue discussed in Chapter 1 of Hjort et al. (2010). A prior distribution represents prior belief, learnt before data is observed, but it also fully specifies the Bayesian learning model. As more data arrives it is expected that the learning model goes in the right direction and here we must add that it does so at a reasonable convergence rate. If it does not, then the learning model (prior) has not been set well, even though the prior might be appropriate as representing prior beliefs. We claim that priors that satisfy Theorem 1 and therefore do not respect the Occam's razor principle will invariably lead to bad learning models. Figure 3 illustrates this point for the case of exponential priors.

4 Cox proportional hazard model with time-varying frailty

In this section we will present an example of a Cox proportional hazard model with time-varying frailty that will be applied to analyse data on chronic granulomatous disease. This model requires the derivation of two new priors. We need to derive a prior for the precision of a Gaussian model component used as a random effect term and we also need a prior for the lag-one correlation parameter in an auto-regressive process of first order (an AR(1) model).

The Cox proportional hazard model is a popular type of survival model that can be fitted to recurrent event data. It assumes that the time-varying hazard for the ith subject can be expressed as

$$h(t; i) = h_0(t) \exp(\eta_i),$$

where the combined risk variable η_i in most cases depends on subject-specific covariates $\boldsymbol{z}_i$. The function $h_0(t)$ is the baseline hazard, see Fleming and Harrington (2005) for further details and applications of model. In the given example, our main focus is on the inclusion of a subject-specific and possibly time-dependent frailty term in η_i, and also on the model for the baseline hazard. Both of these model components will benefit from using our new priors.

4.1 Subject-specific term: independent Gaussian random effects

The notion of frailty is similar to the classical notion of a random effect, and it provides a convenient way to introduce association and unobserved heterogeneity in these models. Using a subject-specific log-normal frailty, we arrive at

$$\eta_i = \boldsymbol{z}_i^T \boldsymbol{\beta} + v_i,$$

where $\boldsymbol{v}$ are i.i.d. Gaussian variables with precision τ. Hence, we need to construct a PC prior on the precision parameter τ for a Gaussian random effect. The natural base model is the absence of random effects, which corresponds to $\tau = \infty$. This base model leads to a useful negative result, which is similar to Theorem 1.

Theorem 2. *Let $\pi_\tau(\tau)$ be a prior for $\tau > 0$ where $E(\tau) < \infty$, then $\pi_d(0) = 0$ and the prior overfits.*

The proof is given in Appendix B.1.

The negative result in Theorem 2 is reflected in the behaviour of the overfitting priors. In particular, priors that overfit induce qualitatively different priors on the random effect v_i. In the following theorem, we derive the behaviour near $v_i = 0$ and $v_i = \infty$ and show that any prior on the distance scale that is non-zero in a

neighbourhood of the base model induces an infinite spike at $v_i = 0$, while priors that overfit induce a bounded prior.

Theorem 3. *A prior on the precision τ that does not overfit induces a marginal prior on the Gaussian random effect v that has at least a weak logarithmic spike at zero. That is, near $v = 0$,*

$$q\pi(v) \geq C \log\left(\frac{\delta^2}{v^2}\right),$$

where $C > 0$ and $\delta > 0$ is any number such that $\pi_d(r) > 0$ for every $r \in [0, \delta]$. Conversely, if $\pi_d(r) = \mathcal{O}(r^t)$ for any $t > 0$, $\pi(v)$ is bounded for all $v \geq 0$.

Theorem 3 is proved in a more general form in Appendix B.2, where the tail behaviour of $\pi(v)$ is linked to the tails of $\pi_d(r)$.

Our new PC prior for τ is derived in Appendix A.2 as a type-2 Gumbel distribution with no integer moments

$$\pi(\tau) = \frac{\lambda}{2}\tau^{-3/2}\exp\left(-\lambda/\sqrt{\tau}\right), \qquad \lambda > 0. \tag{3}$$

This is the generic new prior for the precision parameter of any model component where the precision matrix is $\tau\boldsymbol{R}$ for a fixed matrix $\boldsymbol{R}$. This also includes the case where $\boldsymbol{R}$ is not of full rank, like in popular intrinsic models, spline and thin-plate spline models (Rue and Held, 2005). Using (approximate) parameter orthogonality, this is also the new prior for the precision of the Student-t distribution which can be combined with the computed prior for the d.o.f. in Section 3.

The parameter λ determines the magnitude of the penalty for deviating from the base model and higher values increase this penalty. As previously, we can determine λ by imposing a notion of *scale* on the random effects. This requires the user to specify (U, α) so that $\text{Prob}(1/\sqrt{\tau} > U) = \alpha$. This implies that $\lambda = -\ln(\alpha)/U$. As a rule of thumb, the marginal standard deviation of $\boldsymbol{v}$, after the type-2 Gumbel distribution for τ is integrated out, is about $0.31U$ when $\alpha = 0.01$. This means that the choice $(U = 0.968, \alpha = 0.01)$ gives $\text{Stdev}(\boldsymbol{v}) \approx 0.3$. Note that thinking about the marginal standard deviation of a random effect is more direct and intuitive than choosing hyperparameters of a given prior. For example, a standard approach here would be to assign a Gamma distribution to the conditional precision τ, which later needs to be integrated out to obtain the distribution of the random effect.

The new prior is displayed in Figure 4 for $(U = 0.3/0.31, \alpha = 0.01)$, together with the popular Gamma$(1, a)$ prior, where the shape is 1 and rate is a. In applications, the "art" is to select an appropriate value for a. We selected a so that the marginal variance for the random effects are equal for the two priors. Panel (a) shows the two priors on the precision scale and panel (b) shows the two priors on the distance scale. The priors for low precisions are quite different, and so are the tail behaviours. For large τ, the new prior behaves like $\tau^{-3/2}$, whereas the Gamma-prior goes like $\exp(-a\tau)$. This is a direct consequence of the importance the new prior gives to the base model, i.e. the absence of random effects. Panel (b) demonstrate that the Gamma-prior has density zero at distance zero, and hence, do not prevent over-fitting.

4.2 Baseline hazard term: Dependent Gaussian random effects

A full Bayesian analysis also requires a model for the baseline hazard. A natural choice is to consider the log baseline hazard as piecewise constant function on small time intervals, and impose smoothness to penalise deviations from a constant, see for example Fahrmeir and Tutz (2001). This is an interesting case in which the traditional specification of the penalty term makes it essentially impossible to prescribe a good prior: sometimes it is our fault that it's difficult to specify priors (Sørbye and Rue, 2014). The difficulty lies in the scaling of the random effect, which is automatically corrected by PC priors.

Let $[0, T]$ be the time interval of interest, and divide that interval into n equidistant (for simplicity) intervals $0 < t_1 < t_2 < \cdots < t_{n-1} < T$. Let h_i, $i = 1, \ldots, n$ denote the log baseline hazard in the ith interval. The

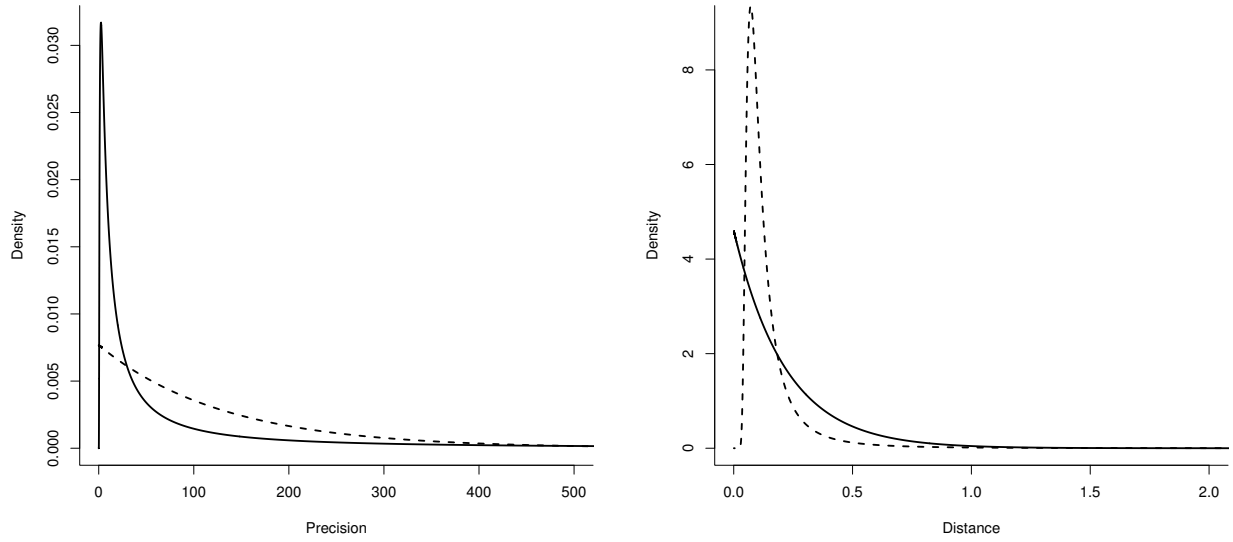

(a) (b)

Figure 4: Panel (a) displays the new prior (solid) with parameters $(U = 0.3/0.31, \alpha = 0.01)$, and the Gamma(shape $= 1$, rate $= a$) prior (dashed). The value of a is computed so that the marginal variance for the random effect are equal for the two priors, which leads to $a = 0.0076$. Panel (b) shows the same two priors on the distance scale demonstrating that the density for the Gamma-prior is zero at distance zero.

first order random walk (RW1) model imposes smoothing penalty among neighbour h_i's,

$$\pi(\boldsymbol{h} \mid \tau) \propto (\tau\tau^*)^{(n-1)/2} \exp\left(-\frac{\tau\tau^*}{2}\sum_{i=2}^{n}(h_i - h_{i-1})^2\right). \tag{4}$$

Here, τ^* is a positive scaling constant that has been added to help alleviate the aforementioned scaling difficulties, while τ denotes the precision which can also be viewed as a smoothing parameter. To separate the baseline hazard from the intercept, we impose the constraint $\sum_i h_i = 0$. The base model is a constant (in time) baseline hazard, which corresponds to infinite smoothing, $\tau = \infty$. With the sum-to-zero constraint the RW1 model is proper but not standardised. More explicitly, the characteristic marginal variance—defined by Sørbye and Rue (2014) to be the geometric mean of the marginal variances—is $\mathcal{O}(n)$. The practical implication of this is that, for a fixed prior on τ, the implied prior on the variance changes with n. This leads to the unpalatable situation where a relatively uninformative prior for one problem can be very informative for another (Sørbye and Rue, 2014). The situation is even worse for smoother models as the characteristic marginal variance of a kth order random walk model grows like $\mathcal{O}(n^k)$.

Bearing this in mind, the scaling constant τ^* in (4) is selected to give a characteristic marginal variance equal to 1 when $\tau = 1$. For the RW1 model above, this implies that $\tau^* \approx 6.63/n$. For more general structure matrices $\boldsymbol{R}$, it is straightforward compute this quantity exactly (Sørbye and Rue, 2014). Using the argument in Section 4.1, the PC prior is once again a type-2 Gumbel prior (3) for the precision τ.

4.3 Model with time-dependent frailty applied to real data

An interesting extension to the subject specific frailty model is to allow the frailty term to depend on time (Yau and McGilchrist, 1998), leading to a time-dependent combined risk variable $\eta_i(t)$. Anticipating a positive correlation in time, it is natural to model this time dependent risk using a continuous time Ornstein-Uhlenbeck process or its discrete time version (the AR(1) model). The stationary AR(1) model for subject i's specific frailty is given by

$$v_{it} \mid \{v_{is}, s < t\} \sim \mathcal{N}(\rho v_{i,t-1}, \tau_v(1-\rho^2)),$$

parameterised so that τ_v is the marginal precision. Here, ρ is the lag-one correlation. In the remainder of this section, we will construct a PC prior for $\rho \in [-1, 1]$. Similar arguments can be used to construct a PC prior for non-negative ρ, however the results were not affected by using the prior on the extended space.

For this model component, the natural base model (keeping the marginal precision constant) is a time-constant frailty, which gives a singular limit of $\rho = 1$. The corresponding PC prior for ρ, derived in Appendix A.3, is

$$\pi(\rho) = \frac{\lambda \exp\left(-\lambda\sqrt{1-\rho}\right)}{\left(1-\exp(-\sqrt{2}\lambda)\right)\ 2\sqrt{1-\rho}}, \tag{5}$$

where $\lambda > 0$ determines the degree of penalisation from the base model. To determine λ we request values (U, α) where Prob$(\rho > U) = \alpha$, see the Appendix for details. For a fixed correlation ρ, the base model is the constant zero which gives the type-2 Gumbel prior in (3) for the precision τ_v .

We end this section by analysing the CGD data (Fleming and Harrington, 2005) available in `R` as the `cgd` dataset in the `survival` package. This data set consists of 128 patients from 13 hospitals with chronic granulomatous disease (CGD). These patients participated in a double-blinded placebo controlled randomised trial, in which a treatment using gamma interferon (γ-IFN) was used to avoid or reduce the number of infections suffered by the patients. The recorded number of CGD infections for each patient ranged from zero to a maximum of seven, and the survival times are given as the times between recurrent infections on each patient. We follow Yau and McGilchrist (1998) and introduce a deterministic time dependent covariate for each patient, given as the time since the first infection (if any). Additionally, we include the covariates treatment (placebo or γ-IFN), inherit (pattern of inheritance), age (in years), height (in cm), weight (in kg), propylac (use of prophylactic antibiotics at study entry), sex, region (US or Europe), and steroids (presence of corticosteroids) (Manda and Meyer, 2005; Yau and McGilchrist, 1998). The covariates age, height and weight were scaled before the analysis. We will use a constant prior for the intercept and independent Gaussian priors with precision 0.001 for the covariate weights.

The computations were performed using the `R-INLA` package, by rewriting the model into a larger Poisson regression, see Fahrmeir and Tutz (2001) for a more general discussion and Martino et al. (2010) for `R-INLA` specific details. The prior specifications are as follows. We used a constant prior for the intercept and independent zero mean Gaussian prior with low precision for all the fixed effects. For the log baseline hazard with $n = 25$ segments, we used the type-2 Gumbel prior with parameters $(U = 0.15/0.31, \alpha = 0.01)$ giving a marginal standard deviation for the log baseline hazard of about 0.15. The time-dependent frailty was given a type 2 Gumbel prior for the precision with parameters $(U = 0.3/0.31, \alpha = 0.01)$ giving a marginal standard deviation of about 0.3. For the derived prior (5) for ρ, we used the parameters $(U = 1/2, \alpha = 0.75)$. The results are show in Figure 5.

Figure 5 (a) shows the prior density (dashed) and posterior density (solid) for the correlation in the AR(1) model for the frailty. The data hardly alters the prior at all, showing that there is not much information in the data available for this parameter, and we cannot conclude anything about the time-varying frailty. This is somewhat contrary to the findings in Manda and Meyer (2005) and Yau and McGilchrist (1998). Figure 5 (b) displays the log baseline hazard, showing an increasing trend (additional to the deterministic time dependent covariate), but the wide point-wise credible bands give no clear evidence for a time-dependent baseline hazard. With the new prior we are more confident that we do not overfit the data using the more flexible model for the log baseline hazard as we do control the amount of deviation and its shrinkage towards it. The given conclusions are robust to changes in the parameter choices (U, α) for the different model components.

5 Disease mapping using the BYM model

In this section we will discuss how our new approach can be applied within disease mapping, in particular using the popular BYM-model (Besag et al., 1991). This example is particularly interesting since we are required to reparameterise the model in order to make the sequence of base-models interpretable and the parameters more orthogonal. The problem of mapping disease incidence is a huge field within public health

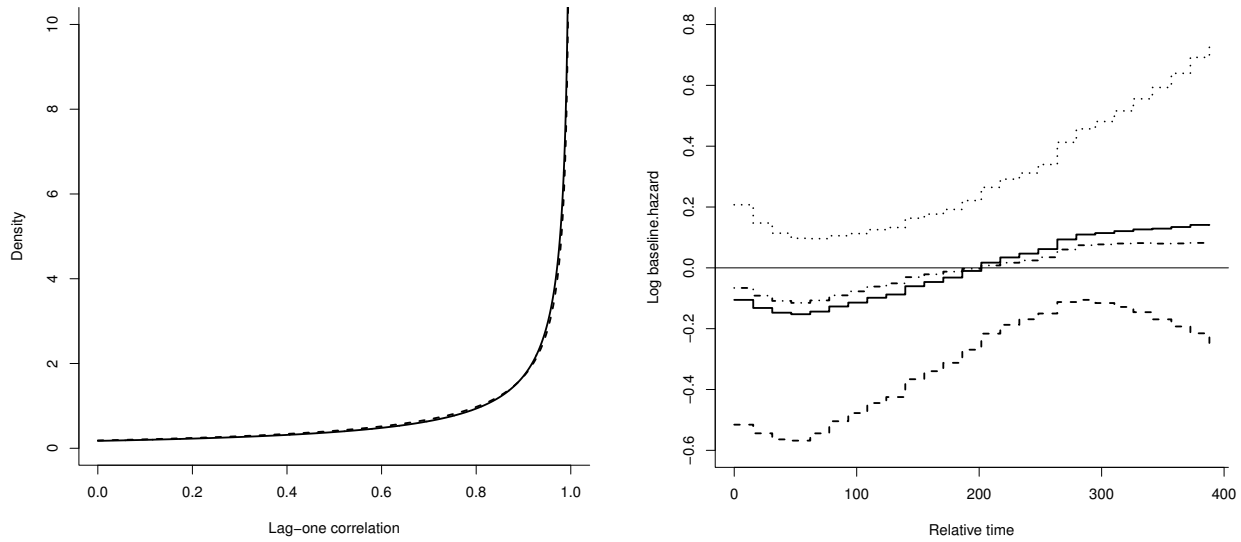

(a) (b)

Figure 5: Panel (a) displays the posterior density (solid) and prior density (dashed) for the lag-one correlation ρ in the AR(1) model for the time-dependent frailty. Panel (b) displays the log baseline hazard, mean (solid), median (dashed-dotted), lower (0.025, dashed) and upper (0.975 , dotted) quantiles.

and epidemiology, and good introductions to the field exist (Lawson, 2006, 2013; Wakefield et al., 2000; Waller and Carlin, 2010).

Disease incidences are usually rare and, in order to anonymise the data, they are commonly available on an aggregated level, e.g. some administrative regions like county or post-code area. The basic model for the observed counts $y_i, i = 1, \ldots, n$, is the Poisson distribution and the observations are assumed to be conditionally independent

$$y_i \mid \boldsymbol{\eta} \sim \text{Poisson}\left(E_i \exp\left(\eta_i\right)\right)$$

where $\{E_i\}$ are the expected number of cases. These are precomputed taking the number of people in the areas, their age distribution and so on, into account. The η_i is the log relative risk, which gives a correction to the expected risk E_i. In the BYM model, the log relative risk is divided into four components

$$\eta_i = \mu + \boldsymbol{z}_i^T \boldsymbol{\beta} + u_i + v_i$$

where μ is the overall intercept, $\boldsymbol{\beta}$ measures the effect of possibly region specific covariates $\boldsymbol{z}_i$, $\boldsymbol{u}$ is a spatial component saying that nearby regions are similar, and $\boldsymbol{v}$ is an independent Gaussian random effect.

The random effect $\boldsymbol{v}$ is a zero mean Gaussian with precision matrix $\tau_v \boldsymbol{I}$. A first order intrinsic Gaussian Markov random field model (Rue and Held, 2005, Ch. 3) was introduced by Besag et al. (1991) as a model for the spatial component $\boldsymbol{u}$. Let $\mathcal{G}$ be the conditional independence graph of $\boldsymbol{u}$, where ∂i denotes the set of neighbours to node i and let $n_{\partial i}$ be the corresponding number of neighbours. The full conditionals are

$$u_i \mid \boldsymbol{u}_{-i} \sim \mathcal{N}\left(\frac{1}{n_{\partial i}} \sum_{j \in \partial i} u_j, n_{\partial i} \tau_u\right)$$

where τ_u is the precision parameter; see Rue and Held (2005, Ch. 3) for details. The full conditionals say that the mean for u_i is the mean of the neighbours, with a precision proportional to the number of neighbours. This model is intrinsic in the sense that the precision matrix has a non-empty null-space. The null-space here is the $\mathbf{1}$-vector, hence the density is invariant to adding a constant to $\boldsymbol{u}$. To prevent confounding with the intercept, we impose the constraint that $\mathbf{1}^T \boldsymbol{u} = 0$, assuming that the graph has only one connected component. To complete the model, we need the prior specification for the intercept and the fixed-effects $\boldsymbol{\beta}$, additional to the prior for the two precision parameters τ_u and τ_v. It is common to assign independent Gamma-priors for τ_u

and τ_v, often of the form $\Gamma(1, a)$ where a_u and a_v are taken to some small values. These choices are usually rather ad-hoc; see for example Bernardinelli et al. (1995) for a discussion.

There are two main issues with the BYM model and the choice of priors. The first is that the spatial component is not *scaled*, which we also discussed for the first order random walk model for the baseline hazard, see Section 4.2. The marginal variance after imposing the $\mathbf{1}^T\boldsymbol{u} = 0$ constraint is not standardised, meaning that any recommended prior (like those suggested by Bernardinelli et al. (1995)) cannot be transferred from one graph to another, since the characteristic marginal variance depends on the graph (Sørbye and Rue, 2014). The second issue is that the structured component $\boldsymbol{u}$ cannot be seen independently from the unstructured component $\boldsymbol{v}$. The independent random effects $\boldsymbol{v}$ are somehow partially included in the spatial random effect $\boldsymbol{u}$ for the case where no spatial dependence is found. Therefore, a proper model should account for this feature in order to avoid identifiability issues. This means that the priors for τ_u and τ_v should be (heavily) dependent, and not independent as it is usually assumed.

In order to resolve these issues, and prepare the model-structure for our new priors, we will reparameterise the model. Let $\boldsymbol{u}^*$ denote a standardised spatial component where the characteristic variance is one. We then rewrite the log relative risk as

$$\eta_i = \mu + \boldsymbol{z}_i^T\boldsymbol{\beta} + \frac{1}{\sqrt{\tau}}\left(\sqrt{1-\phi}\, v_i + \sqrt{\phi}\, u_i^*\right),$$

where $0 \leq \phi \leq 1$ is a mixing parameter and the precision τ controls the marginal precision. The marginal precision contribution from $\boldsymbol{u}^*$ and $\boldsymbol{v}$ is $1/\tau$, whereas the fraction of this variance explained by the spatial term $\boldsymbol{u}^*$ and the random effects $\boldsymbol{v}$, are ϕ and $1-\phi$, respectively. Note that the two hyperparameters (τ, ϕ) are almost orthogonal and naturally allow for independent prior specification.

With this reparameterisation, we can apply our new approach to construct priors. First, we notice that the type-2 Gumbel prior applies to the precision τ, as the natural base model is no effect from $\boldsymbol{u}^*$ and $\boldsymbol{v}$. For a fixed marginal precision, the base model is no spatial dependency i.e. $\phi = 0$. An increased value of ϕ will blend in spatial dependency keeping the marginal precision constant, hence more of the variability will be explained by $\boldsymbol{u}^*$ and the ratio is ϕ. The prior for ϕ is derived in the Appendix A.4 and depends on the graph $\mathcal{G}$. Our notion of *scale* can be used to set (U, α) so that $\text{Prob}(\phi < U) = \alpha$ which determines the degree of penalisation, see the Appendix A.4 for details. This new reparameterisation should shrink towards both no-effect and no-spatial effect and should therefore prevent over-fitting of the data due to unfortunate *ad hoc* prior choices. We will now apply this new parameterisation and priors to two data sets that have been analysed previously, each with its own difficulty and challenges.

5.1 The Sardinia data: no effects

First, we apply the BYM model with the PC priors to insulin-dependent diabetes mellitus (IDDM) counts in Sardinia, which has been analysed by Bernardinelli et al. (1997) and Knorr-Held and Rue (2002). The characteristic feature of this dataset is that it has low counts; There is a total of 619 cases in the 366 districts and there are no covariates. Spatial smoothing is essential to reduce the large variance in the estimates, and to get a realistic picture of the underlying risk surface. The challenge is that Held (2004) has demonstrated that a constant risk is well supported by the posterior distribution, a claim we also have verified using the approach in Sørbye and Rue (2011). With this background, we would expect shrinkage towards this constant when applying the new priors.

We use a constant prior for the intercept, parameters $(U = 0.2/0.31, \alpha = 0.01)$ for the precision τ, which gives a marginal standard deviation of 0.2, and parameters $(U = 1/2, \alpha = 2/3)$ for the prior for ϕ, which gives a $2/3$ probability that the fraction of the marginal variance explained by the random effect $\boldsymbol{v}$ is larger than $1/2$. The results are shown in Figure 6 and Figure 7.

Figure 6 (a) shows the new prior for ϕ (dashed) and the posterior (solid). The new prior is decaying for increasing ϕ until it is close to 1 for which near all the variance is explained by the intrinsic component $\boldsymbol{u}^*$. The posterior for ϕ is mostly uniform giving essentially no information on which component is describing the variance. The posterior for the precision τ is shown in (b), with a median of 46 giving a marginal standard

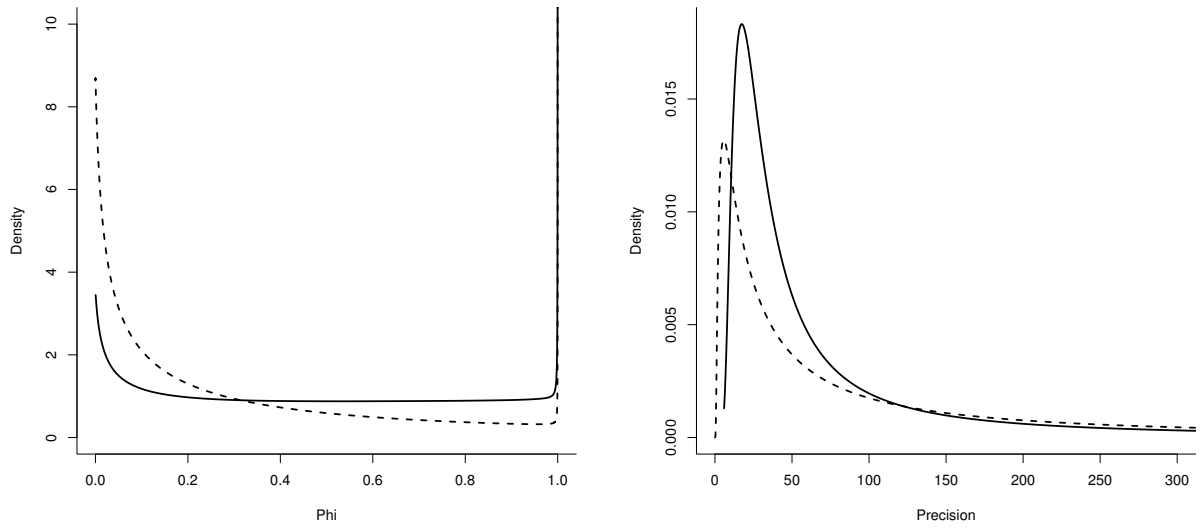

(a) (b)

Figure 6: The results for the Sardinia data using the PC priors and parameterisation of the BYM model. Panel (a) shows the prior density (dashed) and posterior density (solid) for ϕ, and panel (b) shows the prior density (dashed) and the posterior density (solid) for the precision τ.

deviation of 0.16 for the deviation from the overall mean. If we simulate data from the model $\eta_i = 0$, then we obtain a marginal standard deviation of about 0.1 showing the inherent level variability in the data and the model. The new priors seem to do well and do not appear to overfit the data. The results seem to support a constant log relative risk as a zero log relative risk is well within all the posterior credible intervals for $\{\eta_i\}$. Figure 7 shows the map of the posterior mean of the relative risk $\{\exp(\eta_i)\}$.

5.2 Larynx data in Germany: spatial effect and the effect of an ecological covariate

In this example, we will reanalyse the larynx cancer mortality for men, registered in 544 districts of Germany from 1986 to 1990 (Natario and Knorr-Held, 2003). The total number of deaths due to larynx cancer was 7283, which gives an average of 13.4 per region. The interesting part of this model is the semi-parametric estimation of the covariate effect of lung cancer mortality rates in the same period. This covariate acts as an ecological covariate (Wakefield and Lyons, 2010) to account for smoking consumption, which is known to be the most important risk factor of the larynx cancer.

As a smooth model for the ecological covariate, Natario and Knorr-Held (2003) used a second order random walk model

$$\pi(\boldsymbol{x} \mid \tau_x) \propto (\tau_x \tau_x^*)^{(m-2)/2} \exp\left(-\frac{\tau_x \tau_x^*}{2} \sum_{i=2}^{m} (x_i - 2x_{i-1} + x_{i-2})^2\right)$$

where m is the length of $\boldsymbol{x}$. This spline model penalises the estimated second order derivatives and its null space is spanned by $\mathbf{1}$ and $(1, 2, \ldots, m)$. Similar to the spatial component in the BYM model, this model component is not standardised and τ_x^* ensures that the characteristic variance is one. The base model is here a straight line which reduces to a linear effect of the ecological covariate, and the type-2 Gumbel distribution is our new prior.

The log relative risk for this example is given by

$$\eta_i = \mu + f(c_i; \tau_x) + \frac{1}{\sqrt{\tau}}\left(\sqrt{1-\phi}\, v_i + \sqrt{\phi}\, u_i^*\right)$$

where $f(c_i; \tau_x)$ refers to the spline model at location c_i, where the ecological covariate $\boldsymbol{c}$ have been converted into the range $1, 2, \ldots, m$ for simplicity. See Rue and Held (2005, Ch. 3) for more details on this spline

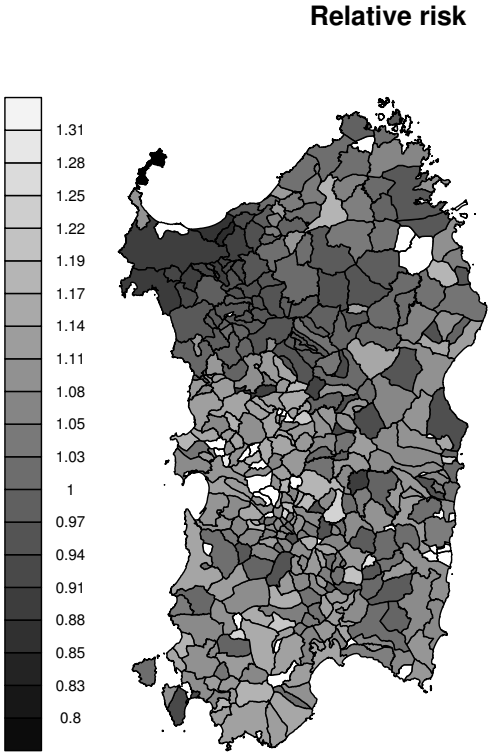

Figure 7: The posterior mean of the relative risk for the Sardinia data.

model, and Lindgren and Rue (2008) for an extension to irregular locations. For $\boldsymbol{u}^*$ and ϕ we used the same parameters as in the previous example, and also for the precision in the spline model we used $(U = 0.2/0.31, \alpha = 0.01)$. Note that choosing prior parameters using the concept of *scale* simply makes prior specification much more transparent. The results are shown in Figure 8.

For the mixing parameter ϕ, the solid line in Figure 8 (a) shows the posterior density while the prior is displayed using dashed lines. Although the prior puts 2/3 probability of $\phi < 1/2$, we notice that the data dominates the prior and makes the posterior concentrate around 1. This implies that only the spatial component contributes to the marginal variance, illustrating that our model (prior) learns well from data. The posterior for the precision τ is shown in panel (b) and it is seen to be more concentrated than in earlier examples due to the relatively high average counts. The effect of the ecological covariate is shown in panel (c) which displays the mean (solid), the lower and upper 0.025-quantiles (dotted) and the best linear fit (blue dashed). The effect of the ecological covariate seems to be shrunk towards the base model, and much more linear than the various estimates by Natario and Knorr-Held (2003) which seems to us to be affected by over-fitting due to their choice of priors. The appropriateness of a linear effect of the ecological covariate, is also verified using the approach in Sørbye and Rue (2011). Due to the design of our new prior, the clear interpretation of the parameters (U, α) and the good learning abilities demonstrated by the new prior in previous examples, we do believe that the high level of shrinking is a data driven effect and not a hidden effect of the prior. Panel (d) shows the map of the posterior mean for the relative risk $\{\exp(\eta_i)\}$.

6 Multivariate PC priors

The examples considered thus far have been essentially univariate, with higher dimensional parameter spaces dealt with using approximate orthogonality. In this section, we will demonstrate that this methodology natu-

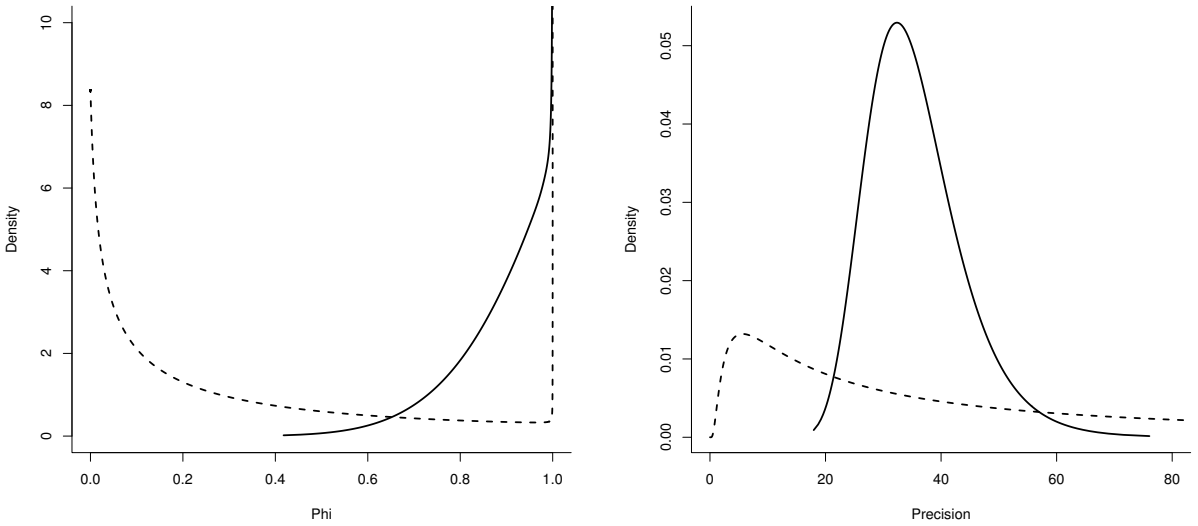

(a)

(b)

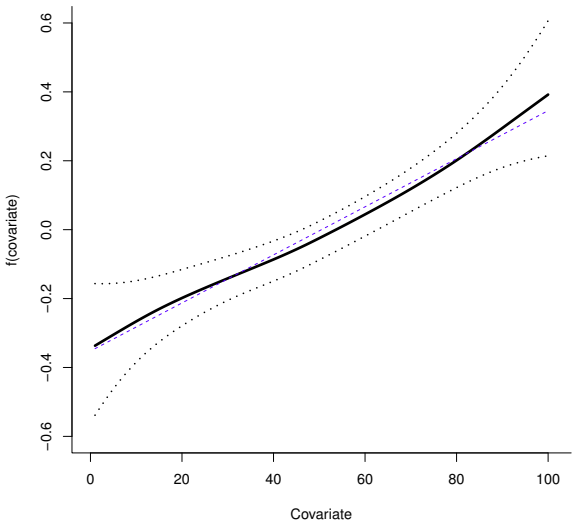

(c)

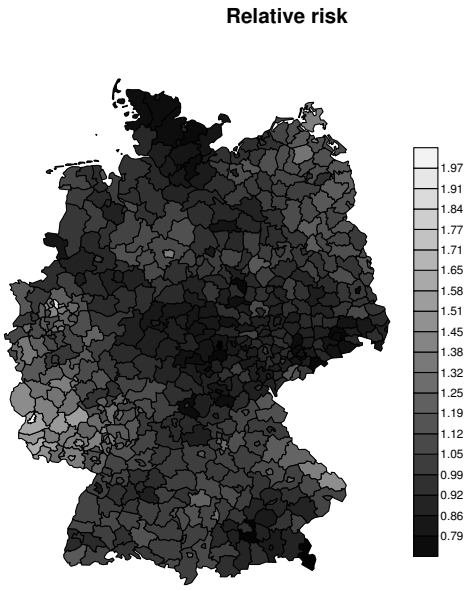

(d)

Figure 8: The results for the larynx data in Germany using the new priors and parameterisation of the BYM model. Panel (a) shows the prior density for ϕ (dashed) and the posterior density (solid). Panel (b) shows the prior density (dashed) and the posterior density (solid) for the precision. Panel (c) shows the effect of the ecological covariate where the solid line is the mean, the dotted lines are the upper and lower 0.025-quantiles, and the blue dashed line is the best linear fit to the mean. Panel (d) shows the posterior mean of the relative risk $\{\exp(\eta_i)\}$.

rally extends to multivariate parameters. This multivariate extension has all the features of the univariate case. We will apply the multivariate PC prior methodology to compute the PC prior for general correlation matrices, and Toeplitz correlation matrices which corresponds to an auto-regressive process.

6.1 Extending the univariate construction

The principles outlined in Section 2.2 can be extended to the multivariate setting $\boldsymbol{\xi} \in \mathcal{M}$ with base model $\boldsymbol{\xi} = \boldsymbol{\xi}_0 \in \mathcal{M}$. As most interesting multivariate parameters spaces are not $\mathbb{R}^n$, we will let $\mathcal{M}$ be a subset of a smooth manifold. For example, when modelling covariance matrices $\mathcal{M}$ will be the manifold of symmetric positive definite matrices, while the set of correlation matrices of a convex subset of that space. A nice introduction to models on manifolds can be found in Byrne and Girolami (2013), where the problem of constructing useful MCMC schemes is also considered.

Assume that $d(\boldsymbol{\xi})$ has a non-vanishing Jacobian. For each $r \geq 0$, the level sets $\boldsymbol{\theta} \in S_r = \{\boldsymbol{\xi} \in \mathcal{M} : d(\boldsymbol{\xi}) = t\}$ are a system of disjoint embedded submanifolds of $\mathcal{M}$ (Lee, 2003, Chapter 8). In the parlance of differential geometry, the submanifolds S_r are the leaves of a foliation and the decomposition $\mathcal{M} = \mathbb{R}_+ \times (\sqcup_{r \geq 0} S_r)$ gives rise to a natural coordinate system on $\mathcal{M}$. Hence the natural lifting of the PC prior concept onto $\mathcal{M}$ is the prior that is exponentially distributed in $d(\boldsymbol{\xi})$ and uniformly distributed on the leaves $S_{d(\boldsymbol{\xi})}$.

In some sense, this above definition is enough to be useful. A simple MCMC or optimisation scheme would proceed in a "coordinate ascent" manner, moving first in the distance direction and then along the leaf S_r. More efficient schemes, however, may be derived from a more explicit density. To this end, we can locally find a mapping $\varphi(\cdot)$ such that

$$(d(\boldsymbol{\xi}), \varphi(\boldsymbol{\xi})) = \boldsymbol{g}(\boldsymbol{\xi}).$$

With this mapping, we get a local representation for the multivariate PC prior as

$$\pi(\boldsymbol{\xi}) = \frac{\lambda}{\left|S_{d(\boldsymbol{\xi})}\right|} \exp\left(-\lambda d(\boldsymbol{\xi})\right) \left|\det(\boldsymbol{J}(\boldsymbol{\xi}))\right|, \tag{6}$$

where $J_{ij} = \frac{\partial g_i}{\partial \xi_j}$ is the Jacobian of $\boldsymbol{g}$. While this definition of multivariate PC priors is significantly more involved than the univariate case, it is still useful. In low dimensions, tools from computational geometry can be used to evaluate (6) approximately. Furthermore, we will show in the next two sections that in two important cases, the structure of the level sets S_r is simple enough to find explicit PC priors for general and Toeplitz correlation matrices.

6.2 A PC prior for covariance matrices

In order to model correlated Gaussian random effects, it is necessary to specify a prior over symmetric positive definite covariance matrices. The most common priors for this purpose are the Wishart and Inverse Wishart distributions, which have convenient conjugacy properties and several extensions to these models that remove some of their basic flaws have been proposed (O'Malley and Zaslavsky, 2008; Huang and Wand, 2013). An alternate method is to note that every covariance matrix can be decomposed as $\boldsymbol{\Sigma} = \boldsymbol{L}\boldsymbol{L}^T$, where $\boldsymbol{L}$ is a lower triangular matrix, and specify a prior directly on the elements of $\boldsymbol{L}$ (Smith and Kohn, 2002). This formulation has the distinct disadvantage that the resulting prior will not be exchangeable and will depend on the ordering of the random effects. In our context, this is not necessarily a problem, as the distributions will often be almost exchangeable, however proper exchangeability can be regained by applying random permutations to the resulting distribution. This has the effect of specifying the prior on *logically* triangular matrices $\boldsymbol{L}$, that is matrices that can be permuted to a lower triangular matrix. In the context of sparsity priors, which are considered by (Smith and Kohn, 2002), constructing a prior directly on $\boldsymbol{L}$ is troublesome as the relationship between the sparsity of $\boldsymbol{\Sigma}$ and the sparsity of $\boldsymbol{L}$ deeply depends on the ordering of the nodes (see Chapter 2 of Rue and Held, 2005, for a detailed discussion).

Consider the PC prior on symmetric positive definite matrices with base model $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_0$.The distance function is given by

$$d(\boldsymbol{\Sigma})^2 = \operatorname{tr}\left(\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma}\right) - p - \log(\det(\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma})) = \operatorname{tr}\left(\boldsymbol{H}(\boldsymbol{\Sigma})\right),$$

where $\boldsymbol{H}(\boldsymbol{\Sigma}) = \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma} - \boldsymbol{I} - \log(\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma})$, which is a symmetric positive semidefinite matrix for all $\boldsymbol{\Sigma}$. Here we have used the identity $\log(\det(\boldsymbol{A})) = \mathrm{tr}(\log(\boldsymbol{A}))$, which holds for any symmetric positive definite matrix $\boldsymbol{A}$. It follows that S_r is the set of symmetric positive semi-definite matrices with trace r^2, which is again a flat space.

As the transformation $\boldsymbol{\Sigma} \to \boldsymbol{H}(\boldsymbol{\Sigma})$ is one-to-one, we suggest using the $\boldsymbol{H}$ parameterisation, in which case the PC prior has density

$$\pi(\boldsymbol{H}) = \lambda \left(\frac{\pi^{1/4p(p-1)} \prod_{k=2}^{p} \Gamma\left(\frac{k+1}{2}\right)}{\Gamma\left(\frac{p(p+1)}{2}\right)} \right) \mathrm{tr}(\boldsymbol{H})^{\frac{1}{4}p(p-1)-\frac{1}{2}} \exp\left(-\lambda\sqrt{\mathrm{tr}(\boldsymbol{H})}\right),$$

where the complicated term in parentheses is the inverse volume of the set of matrices with unit trace (Mittelbach et al., 2012). The PC prior expressed in terms of Σ is given in Appendix A.5. From a computational point of view, the PC prior requires an eigendecomposition of $\boldsymbol{H}$ (or $\boldsymbol{\Sigma}$) and therefore scales cubically in the number of correlated random effects.

A potentially more computationally efficient method would be to fix an ordering of the random effects and use the Cholesky factorisation $\boldsymbol{\Sigma} = \boldsymbol{L}\boldsymbol{L}^T$. In this case, the distance function is $d(\boldsymbol{L})^2 = \|\tilde{\boldsymbol{L}}\|_F^2 - p - 2\sum_i \log(\tilde{L}_{ii})$, where $\tilde{\boldsymbol{L}} = \boldsymbol{L}_0^{-1}\boldsymbol{L}$ and $\boldsymbol{\Sigma}_0 = \boldsymbol{L}_0^T\boldsymbol{L}_0$. Under this parameterisation the leaves S_r are much more complicated and we have been unable to compute the corresponding surface measure.

6.3 A prior for general correlation matrices

Along a different tack, Barnard et al. (2000) parameterise the covariance as $\boldsymbol{\Sigma} = \boldsymbol{D}^{1/2}\boldsymbol{R}\boldsymbol{D}^{1/2}$, where $\boldsymbol{D}$ is a diagonal matrix of marginal variances and $\boldsymbol{R}$ is a correlation matrix, that is a symmetric positive definite matrix with ones on the diagonal. This separation allows us to independently specify priors on the variance of each random effect and the correlation between the random effects. In this section, we show that we can construct a PC prior for the correlation matrix $\boldsymbol{R}$, where the correlation matrix in the base model is the identity matrix, which corresponds to independent components. PC priors corresponding to other base models can be constructed by a linear transformation of corresponding model component.

Consider the model component $\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}^{-1})$ where $\boldsymbol{R}$ is a $q \times q$ correlation matrix. The distance function to the base model is given by $d(\boldsymbol{R}) = \sqrt{-\log(\det(\boldsymbol{R}))}$. This distance function can be greatly simplified by considering a different parameterisation of the set of correlation matrices. Rapisarda et al. (2007) show that every correlation matrix can be written as $\boldsymbol{R} = \boldsymbol{B}\boldsymbol{B}^T$, where $\boldsymbol{B}$ is a lower triangular matrix with first row given by a 1 on the diagonal an zeros in every other position and, for rows $i \geq 2$, entries are given by

$$B_{ij} = \begin{cases} \cos(\theta_{ij}), & j = 1; \\ \cos(\theta_{ij}) \prod_{k=1}^{j-1} \sin(\theta_{ik}), & 2 \leq j \leq i-1; \\ \prod_{k=1}^{i-1} \sin(\theta_{ik}), & j = i; \\ 0, & i+1 \leq j \leq q, \end{cases}$$

where $\theta_{ij} \in [0, \pi]$. The advantage of this parameterisation is that the distance function is now given by $d(\boldsymbol{R}) = \sqrt{2\sum_{i=2}^{q}\sum_{j=1}^{i-1}\gamma_{ij}}$, where $\gamma_{ij} = -\log(\sin(\theta_{ij})) \in [0, \infty)$ are the $p = q(q-1)/2$ parameters. This greatly simplifies the structure of the PC prior as the set S_r is just the simplex $r^2/2 = \sum_{ij}\gamma_{ij}$, which is flat and has volume

$$|S_r| = \frac{\sqrt{p+1}}{p!2^{p/2+1}} r^{2p} = c_p r^{2p}.$$

The determinant of the Jacobian is r^{-1}, so the PC prior becomes

$$\pi(\boldsymbol{\gamma}) = \frac{\lambda p! 2^{p/2+1}}{\sqrt{p+1}} d(\boldsymbol{\gamma})^{-(2p+1)} \exp\left(-\lambda d(\boldsymbol{\gamma})\right). \tag{7}$$

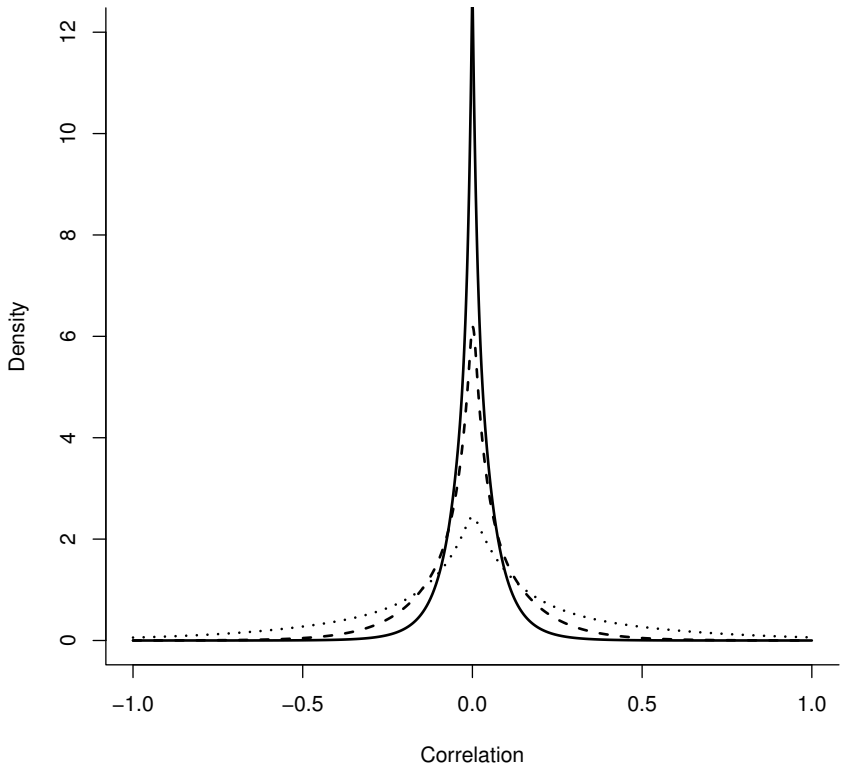

Figure 9: The symmetric marginal prior density for the correlation computed from the PC prior for a 3×3 correlation matrix with $\lambda = 10$ (solid), 5 (dashed) and 2 (dotted).

As the inverse of the mapping $\gamma_{ij}(\theta_{ij}) = \log(\sin(\theta_{ij}))$ has two branches with equal Jacobian, we can compute the prior in $\boldsymbol{\theta}$–space as

$$\pi(\boldsymbol{\theta}) = \frac{\lambda p! 2^{3p/2+1}}{\sqrt{p+1}} d(\boldsymbol{\theta})^{-(2p+1)} \exp\left(-\lambda d(\boldsymbol{\theta})\right) \prod_{i,j} \left| \frac{\cos(\theta_{ij})}{\sin(\theta_{ij})} \right|.$$

The scaling parameter λ controls the degree of compression of the parallelotope with vertices given by the column vectors of $\boldsymbol{R}$. For large values of λ, most of the mass will be concentrated near parallelotopes with unit volume, while for small λ, the volume could be significantly less than one. This parameter may be difficult to visualise in practice, and we suggest calibrating the prior by drawing samples from the model component and selecting a value of λ for which this component behaves appropriately.

There are two complications in interpreting (7). The first, as mentioned in the previous section, is that the corresponding prior is not exchangeable due to the dependence on the Cholesky decomposition on the ordering of the random effects. This can be rectified by summing over all orderings, however we have observed that the pairwise distributions are very similar even without doing this. While we don't necessarily recommend summing out the permutations in practice, for the figures in this section, we have computed the exchangeable PC prior.

The construction of the PC prior can be used to draw samples from $\pi(\boldsymbol{\gamma})$; First sample $r \sim \text{Exp}(\lambda)$, then sample a point uniformly from the simplex S_r (by sampling independent from $\text{Exp}(1)$ and then scale them to get the desired sum), and then transform from back to $\boldsymbol{R}$, randomly selecting the branch of the inverse mapping. Finally, if exchangeability is desired, randomly permute $\boldsymbol{R}$. Figure 9 shows PC prior marginal for one of these correlations for an exchangeable PC prior on a 3×3 correlation matrix, using $\lambda = 10, 5$ and 2. Decreasing λ makes the marginal less centred around zero.

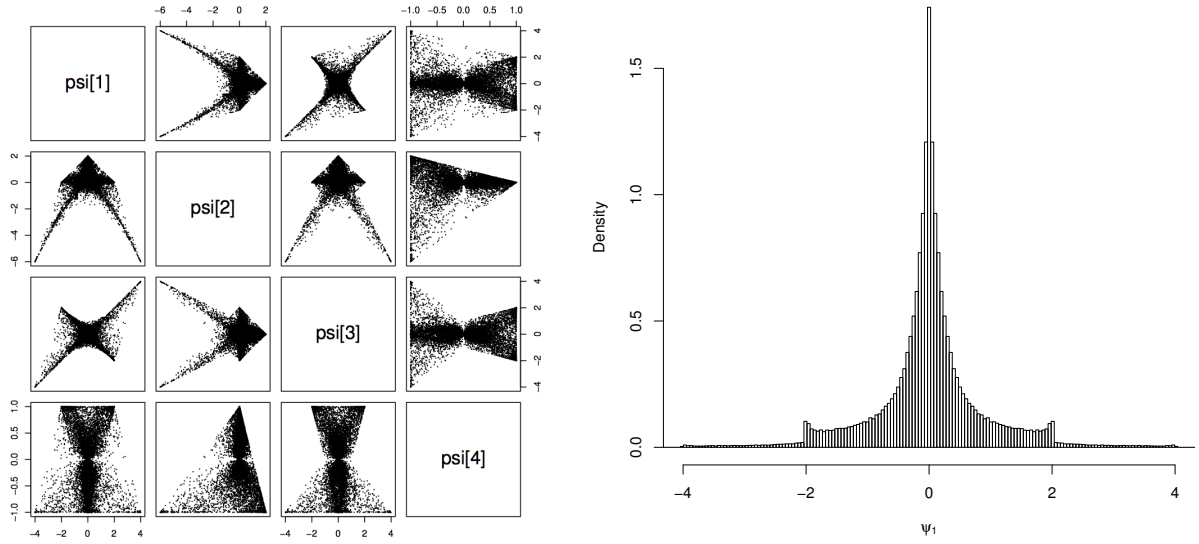

(a) (b)

Figure 10: Panel (a) displays a pairs-plot of 10 000 samples from the prior for $q = 4$ and $\lambda = 0.3$. Panel (b) shows a histogram of the marginal distribution for ψ_1 using 100 000 samples from the same prior.

6.4 A prior for Toeplitz correlation matrices

A related problem to computing the PC prior for a correlation matrix is to compute the PC prior for the coefficients of an AR(q) process,

$$x_t = \psi_1 x_{t-1} + \cdots + \psi_q x_{t-q} + \epsilon_t, \qquad \epsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0, \tau). \tag{8}$$

The corresponding $q \times q$ correlation matrix $\boldsymbol{R}$ is now Toeplitz (Gray, 2002), $R_{ij} = s_{|i-j|}$ where $s_0 = 1$. Although (8) is a natural parameterisation for known $\boldsymbol{\psi}$, it is an awkward parameterisation when these are unknown since the positive definitness requirement makes the space of valid $\boldsymbol{\psi}$ extremely complicated for $q > 3$.

As for general correlation matrices, we can use the invariance property of the PC prior to define the prior for $\boldsymbol{\psi}$ implicitly. The basic idea, which is commonly used when estimating AR(q) parameters, is that the partial correlations $\boldsymbol{\phi} \in [-1, 1]^p$,where $p = q - 1$, are a useful, unconstrained set of parameters for this problem. Furthermore, there is a smooth bijective mappings between the partial correlations and the autocorrelations $\boldsymbol{s}$ given by the Levinson-Durbin recursions (Monahan, 1984; Golub and van Loan, 1996). Expressed in terms of the partial correlations, the distance function is $d(\boldsymbol{\phi})^2 = \sum_i \log(1 - \phi_i^2)$. Similar to the case with general correlation matrices, we can define $\gamma_i = -\log(1 - \phi_i^2)$, for $i = 1, \ldots, p$, and the PC prior follows from (7) as

$$\pi(\boldsymbol{\phi}) = \frac{\lambda p! 2^{5p/2+1}}{\sqrt{p+1}} d(\boldsymbol{\phi})^{-(2p+1)} \exp\left(-\lambda d(\boldsymbol{\phi})\right) \prod_i \left| \frac{\phi_i}{1 - \phi_i^2} \right|. \tag{9}$$

The implied pairwise densities for the AR(4) parameters computed from 10 000 samples with $\lambda = 0.3$ are shown in Figure 10 (a). It is clear that the parameter space has a complicated geometry and that the standard prior, which truncates a multivariate normal prior onto this space (Chib, 1993; Prado and West, 2010), is suboptimal. It is also clear from the geometry of the parameter space that standard MCMC proposals would have difficulty traversing $\boldsymbol{\phi}$–space, whereas reparameterisation in terms of partial correlations flattens the space considerably. Panel (b) shows the histogram for ψ_1 based on 100 000 samples from the same prior. While this marginal is difficult to interpret on its own, it can be understood precisely in terms of the principles underlying the PC prior framework.

7 Discussion

Prior selection is *the* fundamental issue in Bayesian statistics. Priors are the Bayesian's greatest tool, but they are also the greatest point for criticism: the arbitrariness of prior selection procedures and the lack of realistic sensitivity analysis (which is addressed in Roos et al. (2013)) are a serious argument against current Bayesian practice. In this paper, we have provided a principled, widely applicable method for specifying priors on parameters that are difficult to directly elicit from expert knowledge. These PC priors can be vague, weakly informative, or strongly informative depending on the way the user tunes an intuitive scaling parameter. The key feature of these priors is that they explicitly lay out the assumptions underlying them and, as such, these assumptions and principles can be directly critiqued.

PC priors are defined on individual components. This distinguishes PC priors, from reference priors, in which the priors depend on the global model structure. This global dependence is required to ensure a proper posterior. However, the modern applied Bayesian is far more likely to approach their modelling using a component-wise and often additive approach. The directed acyclic graph–approach pioneered by the WinBUGS inference engine, is now a standard tool for specification of general Bayesian models. The additive approach pioneered by Hastie and Tibshirani (1990) is now a standard approach within generalised regression based models. Hence, the ability to specify priors in a componentwise manner is a useful feature.

PC priors enforce shrinkage towards a base model. This shrinkage, however, is not very strong. The aim is to *allow* the posterior to concentrate on the base rather than to force it to. Furthermore, the scale parameter allows the user to specify a reasonable range of parameter values and thus prevent "under-fitting". In the case where the PC prior is placed on a variance parameters, this is conceptually different to the types of Bayesian variable selection priors found in the literature, in which the prior mass is concentrated on sparse sets of parameters (Castillo et al., 2012; Bhattacharya et al., 2012). While these priors are very useful in a number of situations, they are poor general-purpose priors, as the sparsity assumption is extremely informative (by design) and it may not make sense for every problem. Considering priors on the distance scale can, however, be a useful way to construct new Bayesian variable selection priors. Recent results of Bhattacharya et al. (2014) suggest that a normal mixture prior, such as those considered in Theorem 3, will perform optimally in a sparse normal means model provided simple conditions on the behaviour of the implied marginal distribution near zero and infinity are satisfied. We believe that the results outlined in Theorem 3 will be useful in constructing new, optimal, variable selection priors.

The current practice of prior specification is, to be honest, not in a good shape. While there has been a strong growth of Bayesian analysis in science, the research field of "practical prior specification" has been left behind. There are few widely applicable guidelines for how this could or should be done in practice. We hope that with this new principled approach to prior construction, we can reduce the number of "cut and paste" prior choices from other similar research articles, and instead use the derived tools in this paper to specify weakely informative or informative priors with a well defined shrinkage. As always, if the user knows of a better prior for their case, then they should use it. However, having a better default proposal for how to construct priors is a significant advantage.

This is not the whole story of PC priors. We still have to work them out on a case by case basis, construct better guidance for chosing the scaling using knowledge of the global model (like the link-function and the likelihood family), and make them the default choice in packages like `R-INLA`. Not all cases are straight forward. The over-dispersion parameter in the negative Binomial distribution, considered as an extention of the Poisson distribution, cannot be separated from the mean. Hence, we cannot compute the PC prior without knowing a typical value for the mean. We also need to get more experience deriving joint priors for two or more parameters, such as a joint prior for the skewness and kurtosis deviation from a Gaussian.

These new PC prior have made a difference to how we do and see Bayesian analysis. We feel much more confident that the priors we are using do not force over-fitting, and the notion of *scale*, which determine the magnitude of the effects, really simplifies the interpretation of the results. The fact that the prior spesification reduces to a notion of *scale*, makes them very easy to interpret and communicate. We feel that PC priors lay out a new route forward towards more sound Bayesian analysis.

Acknowledgement

The authors acknowledge Simon Barthelmé, Gianluca Biao, Alex Lenkoski, Christian P. Robert, Malgorzata Roos for stimulating discussions and comments related to this work.

A Derivation of PC priors

A.1 The "distance" to a singular model

As PC priors are defined relative to a base model, which is essentially a distinguished point in the parameter space, we occasionally run into the difficulty that this point (denoted $\xi = 0$) is fundamentally different from the points $\xi > 0$. In particular, the base model is occasionally singular to the other distributions and we need to define a useful notion of distance from a singular point in the parameter space. We are saved in the context of this paper by noting that the singular model $\pi(x|\xi = 0)$ is the end point of a curve in model space $t \to \pi(x|\xi = t)$, $t \geq 0$.

Consider the ϵ-distance $d_\epsilon(t) = \sqrt{2\text{KLD}\left(\pi(x|\xi = t) \| \pi(x|\xi = \epsilon)\right)}$, which is finite for every $\epsilon > 0$. If $d_\epsilon(t) = \mathcal{O}(\epsilon^{-p})$, then we can define the renormalised distance $\tilde{d}_\epsilon(t) = \epsilon^{p/2} d_\epsilon(t)$. Using this renormalisation, $\tilde{d}(t) = \lim_{t \downarrow 0} \tilde{d}_\epsilon(t)$ is finite and parameterisation invariant and we can use it to define PC priors for singular models.

A.2 The PC prior for the precision in a multivariate Normal distribution

Let $\mathcal{N}^{(p)}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denote a multivariate normal distribution with dimension p. The Kullback-Leibler divergence from $\mathcal{N}_1^{(p)}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ to $\mathcal{N}_0^{(p)}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ is

$$\text{KLD}\left(\mathcal{N}_1^{(p)} \| \mathcal{N}_0^{(p)}\right) = \frac{1}{2}\left\{\text{tr}\left(\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma}_1\right) + (\boldsymbol{\mu}_0 - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_0^{-1}(\boldsymbol{\mu}_0 - \boldsymbol{\mu}_1) - p - \ln\left(\frac{|\boldsymbol{\Sigma}_1|}{|\boldsymbol{\Sigma}_0|}\right)\right\}.$$

In our setting, $\mathcal{N}_1^{(p)}$ denotes the flexible model and $\mathcal{N}_0^{(p)}$ the base model.

To derive the PC prior for τ where $\boldsymbol{\Sigma}_1 = \boldsymbol{R}/\tau$, $\boldsymbol{R}$ is a fixed matrix and $\boldsymbol{\Sigma}_0 = \mathbf{0}$, we will study the limiting behaviour when $\boldsymbol{\Sigma}_0 = \boldsymbol{R}/\tau_0$ for a high fixed value of τ_0. In the end we will look at the limit $\tau_0 \to \infty$. For simplicity, assume $\boldsymbol{R}$ has full rank. We then get

$$\text{KLD} = \frac{p}{2}\frac{\tau_0}{\tau}\left(1 + \frac{\tau}{\tau_0}\ln\left(\frac{\tau}{\tau_0}\right) - \frac{\tau}{\tau_0}\right) \longrightarrow \frac{p}{2}\frac{\tau_0}{\tau}, \qquad \tau \ll \tau_0,$$

and $d(\tau) = \sqrt{p\tau_0/\tau}$. With an exponential prior for d with rate $\lambda = \theta/\sqrt{p\tau_0}$ we get

$$\pi(\tau) = \frac{\theta}{2}\tau^{-3/2}\exp\left(-\theta/\sqrt{\tau}\right),$$

which is a type-2 Gumbel distribution. To infer θ from a notion of *scale*, we request (U, α) saying that large standard deviations are less likely,

$$\text{Prob}\left(\frac{1}{\sqrt{\tau}} > U\right) = \alpha.$$

This gives $\theta = -\ln(\alpha)/U$. We can now take the limit $\tau_0 \to \infty$ by choosing λ so that θ is kept constant. In the case where $\boldsymbol{R}$ is not of full rank, we use the generalised inverse and determinant, and we arrive at the same result.

A.3 The PC prior for the lag-one correlation in an AR(1) model

Depending on the application, we will use either $\rho = 0$ (no dependence in time) or $\rho = 1$ (no changes in time) as the base-model. We assume the AR(1) model has unit marginal precision. Consider first the case where $\rho = 0$ represents the base model and let p be the dimension. In this case $\boldsymbol{\Sigma}_0 = \boldsymbol{I}$, $\boldsymbol{\Sigma}_1 = \left(\rho^{|i-j|}\right)$, and then it follows that $2\text{KLD}(\rho) = (1-p)\log(1-\rho^2)$ and $d(\rho) = \sqrt{2\text{KLD}(\rho)}$. With an exponential prior for $d(\rho)$ with rate $\lambda = \theta/\sqrt{p-1}$, we get

$$\pi(\rho) = \frac{\theta}{2}\exp\left(-\theta\sqrt{-\ln(1-\rho^2)}\right)\ \frac{|\rho|}{(1-\rho^2)\sqrt{-\ln(1-\rho^2)}}, \qquad |\rho| < 1.$$

To infer θ we request (U, α) saying that large absolute correlation are less likely $\text{Prob}(|\rho| > U) = \alpha$. This gives $\theta = -\ln(\alpha)/\sqrt{-\ln(1-U^2)}$.

For the case where $\rho = 1$ is the base-model, we need to do a limiting argument, as the base-model is singular. Let $\boldsymbol{\Sigma}_1 = \left(\rho^{|i-j|}\right)$ and $\boldsymbol{\Sigma}_0 = \left(\rho_0^{|i-j|}\right)$ where ρ_0 is close to 1 and $\rho < \rho_0$. We then find that $d(\rho) = C\sqrt{1-\rho}$ for some constant C not depending on ρ. For simplicity we take $C = 1$. Since $0 \le d(\rho) \le \sqrt{2}$, we take a truncated exponential distribution for $d(\rho)$ with rate θ and we get

$$\pi(\rho) = \frac{\theta\exp\left(-\theta\sqrt{1-\rho}\right)}{\left(1-\exp\left(-\sqrt{2}\theta\right)\right)\ 2\sqrt{1-\rho}}.$$

To infer θ we request (U, α) saying that $\text{Prob}(\rho > U) = \alpha$, which is given implicitly by

$$\frac{1-\exp\left(-\theta\sqrt{1-U}\right)}{1-\exp\left(-\sqrt{2}\theta\right)} = \alpha$$

provided that α is larger than the lower limit $\sqrt{(1-U)/2}$. We can now take the limit as $\rho_0 \to 1$ which involves the full expression for C, and as it turns out that the solution is the same.

A.4 The PC prior for the mixing parameter in the BYM-model

We will now derive the PC prior for the mixing parameter ϕ in the new parameterisation for the BYM model. Let $\boldsymbol{u}$ be a n-dimensional standardised Gaussian model with zero mean and precision matrix $\boldsymbol{R} > 0$, $\boldsymbol{v}$ be an independent zero mean random effects with unit variance $\mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$, and where the mixing parameter ϕ satisfies $0 \le \phi \le 1$. The more flexible model is $\sqrt{1-\phi}\ \boldsymbol{v} + \sqrt{\phi}\ \boldsymbol{u}$, and the base-model is $\boldsymbol{v}$ (i.e. the model flexible model when $\phi = 0$). Let $\boldsymbol{\Sigma}_0 = \boldsymbol{I}$ and $\boldsymbol{\Sigma}_1(\phi) = (1-\phi)\boldsymbol{I} + \phi\boldsymbol{R}^{-1}$, then

$$\begin{aligned} 2\,\text{KLD}(\phi) &= \text{tr}(\boldsymbol{\Sigma}_1(\phi)) - n - \ln|\boldsymbol{\Sigma}_1(\phi)| \\ &= n\phi\left(\frac{1}{n}\text{tr}(\boldsymbol{R}^{-1}) - 1\right) - \ln|(1-\phi)\boldsymbol{I} + \phi\boldsymbol{R}^{-1}| \end{aligned}$$

and $d(\phi) = \sqrt{2\text{KLD}(\phi)}$. The interesting case is when $\boldsymbol{R}$ is sparse, for which $\text{tr}(\boldsymbol{R}^{-1})$ is quick to compute (Erisman and Tinney, 1975; Rue and Martino, 2007). For the determinant term, we can massage the expression to facility the speedup of computing with sparse matrices. Using the matrix identity $(\boldsymbol{I} + \boldsymbol{A}^{-1})^{-1} = \boldsymbol{A}(\boldsymbol{A} + \boldsymbol{I})^{-1}$, we get

$$|(1-\phi)\boldsymbol{I} + \phi\boldsymbol{R}^{-1}| = \frac{|\frac{1-\phi}{\phi}\boldsymbol{R} + \boldsymbol{I}|}{|\frac{1}{\phi}\boldsymbol{R}|}.$$

An alternative approach, is to compute the eigenvalues $\{\gamma_i\}$ of $\boldsymbol{R}$, which we need to do only once. Let $\tilde{\gamma}_i = 1/\gamma_i$, and we get

$$|(1-\phi)\boldsymbol{I} + \phi\boldsymbol{R}^{-1}| = \prod_{i=1}^{n}(1-\phi+\phi\tilde{\gamma}_i)\,.$$

In the case where $\boldsymbol{R}$ is singular we introduce linear constraint(s) to ensure that any realisation of $\boldsymbol{u}$ is in its null-space. It is now easier to use the latter computational strategy, but redefine $\tilde{\gamma}_i$ as $1/\gamma_i$ if $\gamma_i > 0$ and $\tilde{\gamma}_i = 0$ if $\gamma_i = 0$.

The truncated exponential prior for $d(\phi)$ gives

$$\pi(\phi) = \frac{\lambda \exp(-\lambda d(\phi))}{1 - \exp(-\lambda d(1))} \left| \frac{\partial d(\phi)}{\partial \phi} \right|$$

where we do not spell out the Jacobian for simplicity. The penalisation parameter λ can be determined by $\text{Prob}(\phi < u) = \alpha$, requiring that $\alpha > d(u)/d(1)$.

A.5 The Jacobian determinant for the Σ parameterisation of the covariance matrix PC prior

Theorem 3.4 and Exercise 3.9 in Higham (2008) show that the differential of the mapping $\boldsymbol{H} = \boldsymbol{\Sigma_0}^{-1}\boldsymbol{\Sigma} - \boldsymbol{I} - \log(\boldsymbol{\Sigma_0}^{-1}\boldsymbol{\Sigma})$ is

$$d\boldsymbol{H} = \frac{1}{2\pi i} \int_{C(\boldsymbol{\Sigma})} (z - 1 - \log(z)) \left(z\boldsymbol{I} - \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma}\right)^{-1} \boldsymbol{\Sigma}_0^{-1} d\boldsymbol{\Sigma} \left(z\boldsymbol{I} - \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma}\right)^{-1} dz,$$

where $C(\boldsymbol{\Sigma})$ is a curve surrounding the eigenvalues of $\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma}$. Let $\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma} = \boldsymbol{X}^{-1}\boldsymbol{\Lambda}\boldsymbol{X}$ be an eigendecomposition. Then, if $\text{vec}(\boldsymbol{H})$ is the columwise stacking of $\boldsymbol{H}$, then

$$\begin{aligned} d\,\text{vec}(\boldsymbol{H}) &= \frac{1}{2\pi i} \int_{C(\boldsymbol{\Sigma})} (z - 1 - \log(z)) \left(z\boldsymbol{I} - \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma}\right)^{-1} \otimes \left(z\boldsymbol{I} - \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Sigma}\right)^{-1} dz (\boldsymbol{\Sigma}_0^{-1} \otimes \boldsymbol{I}) d\,\text{vec}(\boldsymbol{\Sigma}) \\ &= (\boldsymbol{X}^{-1} \otimes \boldsymbol{X}^{-1}) \frac{1}{2\pi i} \int_{C(\boldsymbol{\Sigma})} f(z) (z\boldsymbol{I} - \boldsymbol{\Lambda})^{-1} \otimes (z\boldsymbol{I} - \boldsymbol{\Lambda})^{-1} dz (\boldsymbol{X} \otimes \boldsymbol{X})(\boldsymbol{\Sigma}_0^{-1} \otimes \boldsymbol{I}) d\,\text{vec}(\boldsymbol{\Sigma}), \end{aligned}$$

where $f(z) = z - 1 - \log(z)$. From this and Cauchy's integral formula, the Jacobian determinant can be computed as

$$\det(\boldsymbol{J}(\Sigma)) = \det(\boldsymbol{\Sigma_0})^{-q} \prod_{i=1}^{q} f'(\lambda_i) \prod_{i \neq i} f[\lambda_i; \lambda_j],$$

where $f[\lambda_i; \lambda_j] = (f(\lambda_i) - f(\lambda_j))/(\lambda_i - \lambda_j)$ is the first divided difference of $f(\cdot)$.

B Proofs

B.1 Proofs of Theorems 1 and 2

The KLD of approximating the unit precision Student-t with dof ν with a standard Gaussian is for large ν,

$$\text{KLD} = \frac{3}{4}\nu^{-2} + \frac{3}{2}\nu^{-3} + \mathcal{O}(\nu^{-4}).$$

Since $\pi_\nu(\nu)$ has a finite first moment it is $o(\nu^{-2})$ as $\nu \to \infty$. Using the fact that $d = \sqrt{2\,\text{KLD}}$, shows that $\pi_d(d) = o(1)$ as $d \to 0$. The proof of Theorem 2 follows along the same lines.

B.2 The effect of PC priors on Gaussian random effects

Theorem 4. *Let $\pi_d(r)$ be a prior on the distance scale for the flexibility parameter of a Gaussian random effect $v \sim \mathcal{N}(0, r^2)$. The induced prior*

$$\pi(v) = \int_0^\infty \frac{1}{2\pi r} \exp\left(-\frac{v^2}{2r^2}\right) \pi_d(r)\, dr$$

has the following properties. Fix $\delta > 0$.

1. *If* $\pi_d(r) \leq Cr^t$ *for all* $r \in [0, \delta]$ *and for some* $C, t > 0$*, then* $\pi(v)$ *is finite at* $v = 0$*.*

2. *If* $\pi_d(r) \in (0, \infty)$ *for every* $r \in [0, \delta]$*, then* $\pi(v)$ *has a weak logarithmic spike at zero, that is*

$$\pi(v) = \mathcal{O}\left[\log\left(1 + \frac{\delta^2}{v^2}\right)\right], \qquad v \to 0.$$

3. *If* $\int_0^\delta \frac{1}{2\pi r} \exp\left(-\frac{v^2}{2r^2}\right) \pi_d(r)\, dr < \infty$*, then*

$$\pi(v) \geq \mathcal{O}\left(v^{-2} \exp\left(-\frac{v^2}{2\delta^2}\right)\right), \qquad |v| \to \infty.$$

4. *If* $\pi_d(r) \leq (\geq) Cr^{-t}$ *for all* $r \in [0, \delta]$ *and for some* $C, t > 0$*, then*

$$\pi(v) \leq (\geq) \mathcal{O}(|v|^{-t}), \qquad v \to 0.$$

5. *If* $\pi_d(r) \leq (\geq) Cr^{-t}$ *for all* $r > \delta$ *and for some* $C, t > 0$*, then*

$$\pi(v) \leq (\geq) \mathcal{O}(|v|^{-t}), \qquad |v| \to \infty.$$

Proof. Let $r = d(\tau)$. For any $\delta > 0$,

$$\pi(v) = \int_0^\delta \frac{1}{2\pi r} \exp\left(-\frac{v^2}{2r^2}\right) \pi_d(r)\, dr + \int_\delta^\infty \frac{1}{2\pi r} \exp\left(-\frac{v^2}{2r^2}\right) \pi_d(r)\, dr = I_1 + I_2.$$

Examining this splitting, we note that I_1 will control the behaviour of $\pi(v)$ near zero, while I_2 will control the tails.

Assuming that $\int_\delta^\infty r^{-1} \pi_d(r)\, dr < \infty$, we can bound I_2 as

$$\frac{1}{2\pi} \exp\left(-\frac{v^2}{2\delta^2}\right) \int_\delta^\infty r^{-1} \pi_d(r)\, dr \leq I_2 \leq \frac{1}{2\pi} \int_\delta^\infty r^{-1} \pi_d(r)\, dr.$$

To prove parts 1 let $\pi_d(r) \leq Cr^t$, $r \in [0, \delta]$ for some $t > 0$. Substituting this into I_1 and computing the resulting integral using Maple, we get

$$I_1 \leq -\frac{C}{2\pi t}\left(2^{-1/2\,t} |v|^t \Gamma\left(1 - 1/2\,t, 1/2\,\frac{v^2}{\delta^2}\right) - \mathrm{e}^{-1/2\,\frac{v^2}{\delta^2}} \delta^t\right) = \mathcal{O}(1),$$

where $\Gamma(a, x) = \int_x^\infty \exp(-t)\, t^{a-1}\, dt$ is the incomplete Gamma function.

To prove parts 2 and 3, we bound I_1 as follows.

$$\left(\inf_{r \in [0,\delta]} \pi_d(r)\right) \int_0^\delta \frac{1}{2\pi r} \exp\left(-\frac{v^2}{2r^2}\right) dr \leq I_1 \leq \left(\sup_{r \in [0,\delta]} \pi_d(r)\right) \int_0^\delta \frac{1}{2\pi r} \exp\left(-\frac{v^2}{2r^2}\right)$$

$$\frac{1}{4\pi}\left(\inf_{r \in [0,\delta]} \pi_d(r)\right) \mathrm{E}_1\left(\frac{v^2}{2\delta^2}\right) \leq I_1 \leq \frac{1}{4\pi}\left(\sup_{r \in [0,\delta]} \pi_d(r)\right) \mathrm{E}_1\left(\frac{v^2}{2\delta^2}\right)$$

$$\frac{1}{8\pi}\left(\inf_{r \in [0,\delta]} \pi_d(r)\right) \exp\left(-\frac{v^2}{2\delta^2}\right) \log\left(1 + \frac{4\delta^2}{v^2}\right) \leq I_1 \leq \frac{1}{4\pi}\left(\sup_{r \in [0,\delta]} \pi_d(r)\right) \exp\left(-\frac{v^2}{2\delta^2}\right) \log\left(1 + \frac{2\delta^2}{v^2}\right),$$

where $\mathrm{E}_1(x) = \int_1^\infty t^{-1} \exp(-tx)\, dt$ and the third line of inequalities follows using standard bounds in the exponential integral (Abramowitz and Stegun, 1972, formula 5.1.20).

Combining the lower and upper bounds, it follows that if $0 < \inf_{r\in[0,\delta]} \pi_d(r) \leq \sup_{r\in[0,\delta]} \pi_d(r) < \infty$, then $\pi(v)$ has a logarithmic spike near zero. Similarly, the lower bounds show that $\pi(v) \geq Cv^{-2} \exp\left(-\frac{v^2}{2\delta^2}\right)$ as $v \to \infty$.

Part 4 follows by considering let $\pi_d(r) = Cr^{-t}$, $r \in [0, \delta]$ for some $t > 0$. Substituting this into I_1 and computing the resulting integral using Maple, we get

$$I_1 = \frac{C}{2\pi t}\left(|v|^{-t}\Gamma\left(1 + 1/2\,t, 1/2\,\frac{v^2}{\delta^2}\right)2^{t/2} - \delta^{-t}\mathrm{e}^{-1/2\,\frac{v^2}{\delta^2}}\right) \sim \mathcal{O}(v^{-t})$$

as $v \to 0$. We note that $I_1 = \mathcal{O}\left(\exp\left(-v^2/(2\delta^2)\right)\right)$ as $|v| \to \infty$.

To prove part 5, let $\pi_d(r) = Cr^{-t}$, $r \in (\delta, \infty)$ for some $t > 0$. Substituting this into I_2, we get

$$I_2 = \frac{C}{8\pi^2}\,2^{1/2\,t}|v|^{-t}\left(\Gamma\left(1/2\,t\right) - \Gamma\left(1/2\,t, 1/2\,\frac{v^2}{\delta^2}\right)\right) = \mathcal{O}(|v|^{-t}),$$

where we used the identity $\Gamma\left(1/2\,t\right) - \Gamma\left(1/2\,t, 1/2\,\frac{v^2}{\delta^2}\right) \to \Gamma\left(1/2\,t\right)$ as $|v| \to \infty$. □

C The expansion for KLD for Student-t

This expansion is valid for $\nu > 9$ with absolute error less than 1.6E–10.

$$\begin{aligned}\mathrm{KLD}(\nu) \;=\; & \frac{3}{4}\,\nu^{-2} + \frac{3}{2}\,\nu^{-3} + \frac{17}{8}\,\nu^{-4} + \frac{29}{10}\,\nu^{-5} + \frac{61}{12}\,\nu^{-6} + \frac{145}{14}\,\nu^{-7} + \frac{273}{16}\,\nu^{-8} + \frac{119}{6}\,\nu^{-9} + \\ & \frac{869}{20}\,\nu^{-10} + \frac{4\,121}{22}\,\nu^{-11} + \frac{6\,169}{24}\,\nu^{-12} - \frac{30\,035}{26}\,\nu^{-13} - \frac{21\,843}{28}\,\nu^{-14} + \\ & \frac{320\,779}{10}\,\nu^{-15} + \frac{995\,105}{32}\,\nu^{-16} - \frac{28\,689\,547}{34}\,\nu^{-17} - \frac{28\,558\,475}{36}\,\nu^{-18} + \\ & \frac{1\,110\,177\,193}{38}\,\nu^{-19} + \frac{1\,110\,701\,481}{40}\,\nu^{-20} + \mathcal{O}\left(\nu^{-21}\right)\end{aligned}$$

A table of values for $\nu \leq 9$ can be found at `https://code.google.com/p/inla/source/browse/inlaprog/src/dbp.c`.

References

Abramowitz, M. and Stegun, I. (1972). *Handbook of Mathematical Functions*. Dover.

Albert, I., Donnet, S., Guihenneuc-Jouyaux, C., Low-Choy, S., Mengersen, K., and Rousseau, J. (2012). Combining expert opinions in prior elicitation. *Bayesian Analysis*, 7(3):503–532.

Barnard, J., McCulloch, R., and Meng, X. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica*, 10(4):1281–1312.

Bayarri, M. J. and García Donato, G. (2008). Generalization of Jeffreys divergence-based priors for Bayesian hypothesis testing. *Journal of the Royal Statistical Society, Series B*, 70(5):981–1003.

Berger, J. O. (2006). The case for objective Bayesian analysis (with discussion). *Bayesian Analysis*, 1(3):385–402.

Berger, J. O., Bernardo, J. M., and Sun, D. (2009a). The formal definition of reference priors. *The Annals of Statistics*, 37(2):905–938.

Berger, J. O., Bernardo, J. M., and Sun, D. (2009b). The formal definition of reference priors. *The Annals of Statistics*, 37(2):905–938.

Bernardinelli, L., Clayton, D., and Montomoli, C. (1995). Bayesian estimates of disease maps: How important are priors? *Statistics in Medicine*, (14):2411–2431.

Bernardinelli, L., Pascutto, C., Best, N. G., and Gilks, W. R. (1997). Disease mapping with errors in covariates. *Statistics in Medicine*, (16):741–752.

Bernardo, J. M. (1979). Reference posterior distributions for Bayesian inference. *Journal of the Royal Statistical Society, Series B*, 41(2):113–147.

Besag, J., York, J., and Mollié, A. (1991). Bayesian image restoration with two applications in spatial statistics (with discussion). *Annals of the Institute of Statistical Mathematics*, 43(1):1–59.

Bhattacharya, A., Pati, D., Pillai, N. S., and Dunson, D. B. (2012). Bayesian shrinkage. *arXiv preprint arXiv:1212.6088*.

Bhattacharya, A., Pati, D., Pillai, N. S., and Dunson, D. B. (2014). Dirichlet-laplace priors for optimal shrinkage. *arXiv preprint arXiv:1401.5398*.

Byrne, S. and Girolami, M. (2013). Geodesic monte carlo on embedded manifolds. *Scandinavian Journal of Statistics*, 40(4):825–845.

Castillo, I., van der Vaart, A., et al. (2012). Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics*, 40(4):2069–2101.

Chib, S. (1993). Bayes regression with autoregressive errors: A gibbs sampling approach. *Journal of Econometrics*, 58(3):275–294.

Chib, S., Nardari, F., and Shephard, N. (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 108(2):281–316.

Consonni, G. and Veronese, P. (2008). Compatibility of prior specifications across linear models. *Statistical Science*, 23(3):332–353.

Cooke, R. M. (1991). *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Envionmental Ethics & Science Policy. Oxford University Press, USA.

Cox, D. R. and Reid, N. (1987). Parameter orthogonality and approximate conditional inference (with discussion). *Journal of the Royal Statistical Society, Series B*, 49(1):1–39.

Erisman, A. M. and Tinney, W. F. (1975). On computing certain elements of the inverse of a sparse matrix. *Communications of the ACM*, 18(3):177–179.

Fahrmeir, L. and Tutz, G. (2001). *Multivariate Statistical Modelling Based on Generalized Linear Models*. Springer-Verlag, Berlin, 2nd edition.

Fleming, T. R. and Harrington, D. P. (2005). *Counting Processes and Suvival Analsis*. Wiley Series in Probability and Statistics (Book 625). Wiley, 2nd edition.

Fong, Y., Rue, H., and Wakefield, J. (2010). Bayesian inference for generalized linear mixed models. *Biometrics*, 11(3):397–412.

Fonseca, T. C. O., Ferreira, M. A. R., and Migon, H. S. (2008). Objective Bayesian analysis for the Student-t regression model. *Biometrika*, 95(2):325–333.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (with discussion). *Bayesian analysis*, 1(3):515–534.

Gelman, A., Jakulin, A., Pittau, M. G., and Su, Y. S. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2(4):1360–1383.

Geweke, J. (2006). Bayesian treatment of the independent Student-t linear model. *Journal of Applied Econometrics*, 8(S1):S19–S40.

Ghosh, M. (2011). Objective priors: An introduction for frequentists (with discussion). *Statistical Science*, 26(2):187–202.

Goldstein, M. (2006). Subjective Bayesian analysis: Principles and practice (with discussion). *Bayesian Analysis*, 1(3):403–420.

Golub, G. H. and van Loan, C. F. (1996). *Matrix Computations*. Johns Hopkins University Press, Baltimore, 3rd edition.

Gray, R. M. (2002). Toeplitz and circulant matrices: A review. Free book available from http://ee.stanford.edu/∼gray, Department of Electrical Engineering, Stanford University.

Hastie, T. J. and Tibshirani, R. J. (1990). *Generalized Additive Models*, volume 43 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London.

Held, L. (2004). Simultanous posterior probability statements from Monte Carlo output. *Journal of Computational and Graphical Statistics*, 13(1):20–35.

Higham, N. J. (2008). *Functions of matrices: theory and computation*. Siam.

Hjort, N. L., Holmes, C., Müller, P., and Walker, S. G. (2010). *Bayesian Nonparametrics*. Cambridge Series in Statistical and Probabilistic Mathematics (Book 28). Cambridge University Press.

Hodges, J. S. (2013). *Richly Parameterized Linear Models: Additive, Time Series, and Spatial Models Using Random Effects*. Texts in Statistical Science. Chapman & Hall/CRC.

Huang, A. and Wand, M. P. (2013). Simple marginally noninformative prior distributions for covariance matrices. *Bayesian Analysis*, 8(2):439–452.

Jacquier, E., Polson, N. G., and Rossi, P. E. (1994). Bayesian analysis of stochastic volatility models. *Journal of Business and Economic statistics*, 20(1):69–87.

Jacquier, E., Polson, N. G., and Rossi, P. E. (2004). Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. *Journal of Econometrics*, 122(1):185–212.

Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 186(1007):453–461.

Kadane, J. B., Dickey, J. M., Winkler, R. L., Smith, W. S., and Peters, S. C. (1980). Interactive elicitation of opinion for a normal linear model. *Journal of the American Statistical Association*, 75(372):845–854.

Kamary, K. (2014). A discussion on Bayesian analysis : Selecting Noninformative Priors. *arXiv preprint arXiv:1402.6257*.

Kass, R. E. and Wasserman, L. (1996). The selection of prior distributions by formal rules. *Journal of the American Statistical Association*, 91(435):1343–1370.

Knorr-Held, L. and Rue, H. (2002). On block updating in Markov random field models for disease mapping. *Scandinavian Journal of Statistics*, 29(4):597–614.

Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, pages 79–86.

Kurowicka, D. and Cooke, R. M. (2006). *Uncertainty Analysis with High Dimensional Dependence Modelling*. Wiley Series in Probability and Statistics. Wiley.

Lange, K., Little, R., and Taylor, J. (1989). Robust statistical modeling using the t distribution. *Journal of the American Statistical Association*, 84(408):881–896.

Lawson, A. B. (2006). *Statistical Methods in Spatial Epidemiology*. Wiley Series in Probability and Statistics (Book 657). Wiley, 2nd edition.

Lawson, A. B. (2013). *Bayesian Disease Mapping: Hierarchical Modeling in Spatial Epidemiology*. Chapman & Hall/CRC Interdisciplinary Statistics. Chapmann & Hall/CRC, 2nd edition.

Lee, J. M. (2003). *Smooth manifolds*. Springer.

Lindgren, F. and Rue, H. (2008). A note on the second order random walk model for irregular locations. *Scandinavian Journal of Statistics*, 35(4):691–700.

Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The SPDE approach (with discussion). *Journal of the Royal Statistical Society, Series B*, 73(4):423–498.

Lunn, D., Spiegelhalter, D., Thomas, A., and Best, N. (2009). The BUGS project: Evolution, critique, and future directions. *Statistics in Medicine*, 28(25):3049–3067.

Manda, S. O. M. and Meyer, R. (2005). Bayesian inference for recurrent events data using time-dependent frailty. *Statistics in Medicine*, 24(8):1263–1274.

Martino, S., Akerkar, R., and Rue, H. (2010). Approximate Bayesian inference for survival models. *Scandinavian Journal of Statistics*, 28(3):514–528.

Martins, T. G. and Rue, H. (2013). Prior for exibility parameters: the Student's t case. Technical report S8-2013, Department of mathematical sciences, NTNU, Norway.

Martins, T. G., Simpson, D., Lindgren, F., and Rue, H. (2013). Bayesian computing with INLA: New features. *Computational Statistics & Data Analysis*, 67(xx):68–83.

Masreliez, C. and Martin, R. (1977). Robust Bayesian estimation for the linear model and robustifying the Kalman filter. *IEEE Transactions on Automatic Control*, 22(3):361–371.

Mittelbach, M., Matthiesen, B., and Jorswieck, E. A. (2012). Sampling uniformly from the set of positive definite matrices with trace constraint. *Signal Processing, IEEE Transactions on*, 60(5):2167–2179.

Monahan, J. F. (1984). A note on enforcing stationarity in autoregressive-moving average models. *Biometrika*, 71(2):403–404.

Natario, I. and Knorr-Held, L. (2003). Non-parametric ecological regression and spatial variation. *Biometrical Journal*, 45:670–688.

O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E., and Rakow, T. (2006). *Uncertain Jedgements: Eliciting Experts' Probabilities*. Statistics in Practice (book 35). Wiley.

O'Malley, A. J. and Zaslavsky, A. M. (2008). Domain-level covariance analysis for multilevel survey data with structured nonresponse. *Journal of the American Statistical Association*, 103(484):1405–1418.

Pinheiro, J., Liu, C., and Wu, Y. (2001). Efficient algorithms for robust estimation in linear mixed-effects models using the multivariate t distribution. *Journal of Computational and Graphical Statistics*, 10(2):249–276.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003). March*, pages 20–22.

Prado, R. and West, M. (2010). *Time Series: Modeling, Computation and Inference*. Text in Statistical Science Series. Chapman & Hall/CRC.

Rapisarda, F., Brigo, D., and Mercurio, F. (2007). Parameterizing correlations: a geometric interpretation. *IMA Journal of Management Mathematics*, 18(1):55–73.

Robert, C. P., Chopin, N., and Rousseau, J. (2009). Harold Jeffreys's Theory of Probability revisited. *Statistical Science*, 24(2):141–254.

Roos, M., Martins, T. G., Held, L., and Rue, H. (2013). Sensitivity analysis for Bayesian hierarchical models. arxiv:1312.4797, University of Zurich.

Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*, volume 104 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London.

Rue, H. and Martino, S. (2007). Approximate Bayesian inference for hierarchical Gaussian Markov random fields models. *Journal of Statistical Planning and Inference*, 137(10):3177–3192. Special Issue: Bayesian Inference for Stochastic Processes.

Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society, Series B*, 71(2):319–392.

Smith, M. and Kohn, R. (2002). Parsimonious covariance matrix estimation for longitudinal data. *Journal of the American Statistical Association*, 97(460).

Sørbye, S. H. and Rue, H. (2011). Simultanous credible bands for latent Gaussian models. *Scandinavian Journal of Statistics*, 38(4):712–725.

Sørbye, S. H. and Rue, H. (2014). Scaling intrinsic Gaussian Markov random field priors in spatial modelling. *Spatial Statistics*, xx(xx):xx–xx. (to appear).

Spiegelhalter, D. J., Thomas, A., Best, N. G., and Gilks, W. R. (1995). BUGS: Bayesian inference using Gibbs sampling. Version 0.50, MRC Biostatistics Unit, Cambridge.

Stan Development Team (2014). Stan: A C++ Library for Probability and Sampling, Version 2.2. http://mc-stan.org/.

Villa, C. and Walker, S. G. (2014). Objective prior for the number of degrees of freedom of a t distribution. *Bayesian Analysis*, 9(1):197–220.

Wakefield, J. and Lyons, H. (2010). Spatial aggregation and the ecological fallacy. In Gelfand, A. E., Diggle, P. J., Fuentes, M., and Guttorp, P., editors, *Handbook of Spatial statistics*, Handbooks for Modern Statistical Methods, chapter 30, pages 541–558. Chapman & Hall/CRC.

Wakefield, J. C., Best, N. G., and Waller, L. A. (2000). Bayesian approaches to disease mapping. In Elliot, P., Wakefield, J. C., Best, N. G., and Briggs, D. J., editors, *Spatial Epidemiology: Methods and Applications*, pages 104–107. Oxford: Oxford University Press.

Waller, L. and Carlin, B. (2010). Disease mapping. In Gelfand, A. E., Diggle, P. J., Fuentes, M., and Guttorp, P., editors, *Handbook of Spatial statistics*, Handbooks for Modern Statistical Methods, chapter 14, pages 217–243. Chapman & Hall/CRC.

Yau, K. K. W. and McGilchrist, C. A. (1998). ML and REML estimation in survival analysis with time dependent correlated frailty. *Statistics in Medicine*, 17(11):1201–1213.